

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 1 025 514 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention
of the grant of the patent:

12.12.2001 Bulletin 2001/50

(21) Application number: **98950192.9**

(22) Date of filing: **27.10.1998**

(51) Int Cl.7: **G06F 17/30**

(86) International application number:
PCT/GB98/03196

(87) International publication number:
WO 99/23578 (14.05.1999 Gazette 1999/19)

(54) MULTIDIMENSIONAL DATA CLUSTERING AND DIMENSION REDUCTION FOR INDEXING AND SEARCHING

DATENGRUPPIERUNG UND DIMENSIONSVERRINGERUNG MEHRDIMENSIONALER DATEN
ZUM INDIZIEREN UND SUCHE

REGROUPEMENT DE DONNEES MULTIDIMENSIONNELLES ET REDUCTION DES DIMENSIONS
POUR L'INDEXATION ET LA RECHERCHE

(84) Designated Contracting States:
DE FR GB IE

(30) Priority: **31.10.1997 US 960540**

(43) Date of publication of application:
09.08.2000 Bulletin 2000/32

(73) Proprietor: **International Business Machines
Corporation**
Armonk, NY 10504 (US)

(72) Inventors:
• **CASTELLI, Vittorio**
Hursley Park, Winchester, Hamps. SO212JN
(GB)
• **LI, Chung-Sheng**
Ossining, NY 10562 (US)
• **THAMASIAN, Alexander**
Pleasantville, NY 10570 (US)

(74) Representative: **Boyce, Conor**
IBM United Kingdom Limited,
Intellectual Property Law,
Hursley Park
Winchester, Hampshire SO21 2JN (GB)

(56) References cited:
US-A- 5 179 643 **US-A- 5 647 058**

- **NG R ET AL: "EVALUATING
MULTI-DIMENSIONAL INDEXING STRUCTURES
FOR IMAGES TRANSFORMED BY PRINCIPAL
COMPONENT ANALYSIS" STORAGE AND
RETRIEVAL FOR STILL IMAGE AND VIDEO
DATABASES 4, SAN JOSE, FEB. 1 - 2, 1996, no.
VOL. 2670, 1 February 1996, pages 50-61,
XP000642562 SETHI I K;JAIN R C (EDS)**
- **"MULTIDIMENSIONAL INDEX STRUCTURE
WITH MULTI-LEVEL ENTRY AND SKIP-LEVEL
SEARCH FOR PARTIALLY-SPECIFIED
QUERIES" IBM TECHNICAL DISCLOSURE
BULLETIN, vol. 39, no. 11, November 1996,
pages 37-39, XP000679812**

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

EP 1 025 514 B1

DescriptionField of the Invention

[0001] The present invention is related to an improved information system. A particular aspect of the present invention is related to generating and searching compact representations of multidimensional data. A more particular aspect of the present invention is related to generating and searching compact index representations of multidimensional data in database systems using associated clustering and dimension reduction information.

Background

[0002] Multidimensional indexing is fundamental to spatial databases, which are widely applicable to Geographic Information Systems (GIS), Online Analytical Processing (OLAP) for decision support using a large data warehouse, and multimedia databases where high-dimensional feature vectors are extracted from image and video data.

[0003] Decision support is rapidly becoming a key technology for business success. Decision support allows a business to deduce useful information, usually referred to as a *data warehouse*, from an operational database. While the operational database maintains state information, the data warehouse typically maintains historical information. Users of data warehouses are generally more interested in identifying trends rather than looking at individual records in isolation. Decision support queries are thus more computationally intensive and make heavy use of aggregation. This can result in long completion delays and unacceptable productivity constraints.

[0004] Some known techniques used to reduce delays are to pre-compute frequently asked queries, or to use sampling techniques, or both. In particular, applying online analytical processing (OLAP) techniques such as data cubes on very large relational databases or data warehouses for decision support has received increasing attention recently (see e.g., Jim Gray, Adam Bosworth, Andrew Layman, and Hamid Pirahesh, "Data Cube: A Relational Aggregation Operator Generalizing Group-By, Cross-Tab, and Sub-Totals", International Conference on Data Engineering, 1996, New Orleans, pp. 152-160) (Gray). Here, users typically view the historical data from data warehouses as multidimensional data cubes. Each cell (or lattice point) in the cube is a view consisting of an aggregation of interests, such as total sales.

[0005] Multidimensional indexes can be used to answer different types of queries, including:

- find record(s) with specified values of the indexed columns (exact search);
- find record(s) that are within $[a1..a2]$, $[b1..b2]$, ..., $[z1..z2]$ where a , b and z represent different dimensions (range search); and
- find the k most similar records to a user-specified template or example (k -nearest neighbor search).

[0006] Multidimensional indexing is also applicable to image mining. An example of an image mining product is that trademarked by IBM under the name MEDIAMINER, which offers two tools: Query by Image Content (QBIC); and IMAGEMINER, for retrieving images by analyzing their content, rather than by searching in a manually created list of associated keywords.

[0007] QBIC suits applications where keywords cannot provide an adequate result, such as in libraries for museums and art galleries; or in online stock photos for Electronic Commerce where visual catalogs let you search on topics, such as wallpapers and fashion, using colors and texture. Image mining applications such as IMAGEMINER let you query a database of images using conceptual queries like "forest scene", "ice", or "cylinder". Image content, such as color, texture, and contour are combined as simple objects that are automatically recognized by the system.

[0008] These simple objects are represented in a knowledge base. This analysis results in a textual description that is then indexed for later retrieval.

[0009] During the execution of a database query, the database search program accesses part of the stored data and part of the indexing structure; the amount of data accessed depends on the type of query and on the data provided by the user, as well as on the efficiency of the indexing algorithm. Large databases are such that the data and at least part of the indexing structure reside on the larger, slower and cheaper part of the memory hierarchy of the computer system, usually consisting of one or more hard disks. During the search process, part of the data and of the indexing structure are loaded in the faster parts of the memory hierarchy, such as the main memory and the one or more levels of cache memory. The faster parts of the memory hierarchy are generally more expensive and thus comprise a smaller percentage of the storage capacity of the memory hierarchy. A program that uses instructions and data that can be completely loaded into the one or more levels of cache memory is faster and more efficient than a process that in addition uses instructions and data that reside in the main memory, which in turn is faster than a program that also

uses instruction and data that reside on the hard disks. Technological limitations are such that the cost of cache and main memory makes it too expensive to build computer systems with enough main memory or cache to completely contain large databases.

[0010] Thus, there is a need for an improved indexing technique that generates indexes of such size that most or all of the index can reside in main memory at any time; and that limits the amount of data to be transferred from the disk to the main memory during the search process. The present invention addresses such a need.

[0011] Several well known spatial indexing techniques, such as R-trees can be used for range and nearest neighbor queries. Descriptions of R-trees can be found, for example, in "R-trees: A Dynamic index structure for spatial searching" by A. Guttman, ACM SIGMOD Conf. on Management of Data, Boston, MA, June, 1994. The efficiency of these techniques, however, deteriorates rapidly as the number of dimensions of the feature space grows, since the search space becomes increasingly sparse. For instance, it is known that methods such as R-Trees are not useful when the number of dimensions is larger than 8, where the usefulness criterion is the time to complete a request compared to the time required by a brute force strategy that completes the request by sequentially scanning every record in the database. The inefficiency of conventional indexing techniques in high dimensional spaces is a consequence of a well-known phenomenon called the "curse of dimensionality" which is described, for instance, in "From Statistics to Neural Networks" NATO ASI Series, vol. 136, Springer-Verlag, 1994, by V. Cherkassky, J. H. Friedman, and H. Wechsles. The relevant consequence of the curse of dimensionality is that clustering the index space into hypercubes is an inefficient method for feature spaces with a higher number of dimensions.

[0012] The following document is also known: NG R ET AL: "EVALUATING MULTIDIMENSIONAL INDEXING STRUCTURES FOR IMAGES TRANSFORMED BY PRINCIPAL COMPONENT ANALYSIS", STORAGE AND RETRIEVAL FOR STILL IMAGE AND VIDEO DATABASES 4, SAN JOSE, FEB. 1-2 1996, vol. 2670, 1 February 1996, pages 50-61. This document is directed to content-based retrieval in image management systems. To tackle the issue of the high dimensionality of image representations, image analysis and dimension-reducing technique is adopted, see page 52, lines 20-21. This document does however not appear to anticipate the combined features of claim 1. In claim 1 a first step of clustering multidimensional data into groups by partitioning is carried out. Only then is a dimensionality reduction index created. These combined steps do not appear to be taught or rendered obvious from document D1 or any other of the available set of prior art documents.

[0013] Because of the inefficiency associated with using existing spatial indexing techniques for indexing a high-dimensional feature space, techniques well known in the art exist to reduce the number of dimensions of a feature space. For example, the dimensionality can be reduced either by variable subset selection (also called feature selection) or by singular value decomposition followed by variable subset selection, as taught, for instance by C. T. Chen, Linear System Theory and Design, Holt, Rinehart and Winston, Appendix E, 1984. Variable subset selection is a well known and active field of study in statistics, and numerous methodologies have been proposed (see e.g., Shibata et al. An Optimal Selection of Regression variables, Biometrika vol. 68, No. 1, 1981, pp. 45-54. These methods are effective in an index generation system only if many of the variables (columns in the database) are highly correlated. This assumption is in general incorrect in real world databases.

[0014] Thus, there is also a need for an improved indexing technique for high-dimensionality data, even in the presence of variables which are not highly correlated. The technique should generate efficient indexes from the viewpoints of memory utilization and search speed. The present invention addresses these needs.

Summary

[0015] In accordance with the aforementioned needs, the present invention is directed to a method for generating compact representations of multidimensional data as claimed in claim 1. The present invention has features for generating searchable multidimensional indexes for databases. The present invention has other features for flexibly generating the indexes and for efficiently performing exact and similarity searches. The present invention has still other features for generating compact indexes which advantageously limit the amount of data transferred from disk to main memory during the search process.

[0016] One example of an application of the present invention is to multidimensional indexing. Multidimensional indexing is fundamental to spatial databases, which are widely applicable to: Geographic Information Systems (GIS); Online Analytical Processing (OLAP) for decision support using a large data warehouse; and image mining products for mining multimedia databases where high-dimensional feature vectors are extracted from image and video data.

[0017] The present embodiment has yet other features for generating and storing a reduced dimensionality index for the reduced dimensionality clusters.

[0018] Depending on the exact spatial indexing technique used within each individual cluster, the target vector can then be retrieved by using the corresponding indexing mechanism. For example, conventional multidimensional spatial indexing techniques, including but not limited to the R-tree can be used for indexing within each cluster. Alternatively, the intra-cluster search mechanism could be a brute force or linear scan if no spatial indexing structure can be utilized.

[0019] In a preferred embodiment, the dimension reduction step is a singular value decomposition, and the index is searched for a matching reduced dimensionality cluster, based on decomposed specified data. An example of the dimensionality reduction information is a transformation matrix (including eigenvalues and eigenvectors) generated by a singular value decomposition and selected eigenvalues of the transformation matrix.

[0020] Another example of a method for generating multidimensional indexes having features of the present invention includes the steps of: creating a representation of the database to be indexed as a set of vectors, where each vector corresponds to a row in the database and the elements of each vector correspond to the values, for the particular row, contained in the columns for which an index must be generated; the set of vectors is then partitioned using a clustering technique into one or more groups (also called clusters) and clustering information associated with the clustering is also generated and stored; a dimensionality reduction technique is then applied separately to each of the groups, to produce a low-dimensional representation of the elements in the cluster as well as with dimensionality reduction information; and for each reduced-dimensionality cluster, an index is generated using a technique that generates efficient indices for the number of dimensions of the cluster.

[0021] According to another feature of the present invention, the method can be applied separately to each of the reduced-dimensionality clusters, and therefore the method becomes recursive. The process of recursively applying both dimension reduction and clustering terminates when the dimensionality can no longer be reduced.

[0022] According to another feature of the present invention, a dimensionality reduction step could be applied to the entire database as a first step in the index generation (before the database partitioning step). During the partitioning (also called clustering) and dimensionality reduction steps, clustering and dimensionality reduction information is generated for use in the search phase.

[0023] According to still another feature of the present invention, the clustering step could be appropriately selected to facilitate the dimensionality reduction step. This can be accomplished, for instance, by means of a clustering method that partitions the space according to the local covariance structure of the data, instead of minimizing a loss derived from a spatially invariant distance function, such as the Euclidean distance.

[0024] The present invention also has features for assessing if other clusters can contain elements that are closer to the specific data than the farthest of the k most similar element retrieved. As is known in the art, clustering information can be used to reconstruct the boundaries of the partitions, and these boundaries can be used to determine if a cluster can contain one of the k nearest neighbors. Those skilled in the art will appreciate that the cluster boundaries are a simple approximation to the structure of the cluster itself, namely, from the mathematical form of the boundary it is not possible to tell whether there are elements of clusters near any given position on the boundary. As an example consider a case where the database contains two spherical clusters of data, and the two clusters are extremely distant from each other. A reasonable boundary for this case would be a hyperplane, perpendicular to the line joining the centroids of the clusters, and equidistant from the centroids. Since the clusters are widely separated, there is no data point near the boundary. In other cases, the boundary could be very close to large number of elements of both clusters. Accordingly, the present invention also has features to determine if a cluster could contain one or more of the k-nearest neighbors of specified data, using a hierarchy of approximations to the actual geometric structure of each cluster.

[0025] In a preferred embodiment, the present invention is embodied as software stored on a program storage device readable by a machine tangibly embodying a program of instructions executable by the machine to perform method steps for generating compact representations of multidimensional data; efficiently performing exact and similarity searches; generating searchable multidimensional indexes for databases; and efficiently performing exact and similarity searches using the indexes.

Brief Description of the Drawings

[0026] These and other features and advantages of the present invention will become apparent from the following detailed description, taken in conjunction with the accompanying drawings, wherein:

Fig. 1 shows an example of a block diagram of a networked client/server system;

Fig. 2 shows an example of the distribution of the data points and intuition for dimension reduction after clustering;

Fig. 3 shows an example of projecting three points in a 3-D space into a 2-D space such that the projection preserves the relative distance between any two points of the three points;

Fig. 4 shows an example of projecting three points in a 3-D space into a 2-D space where the rank of the relative distance is affected by the projection;

Fig. 5 shows an example of a computation of the distance between points in the original space and the projected

subspace;

Fig. 6 shows an example of a logic flow for generating the multidimensional indexing from the data in the database;

Fig. 7 shows an example of a logic flow for performing dimensionality reduction of the data;

Fig. 8 shows an example of logic flow for an exact search using the index generated without recursive decomposition and clustering;

Fig. 9 shows an example of logic flow for an exact search using the index generated with recursive decomposition and clustering;

Fig. 10 shows an example of logic flow for a k-nearest neighbor search using the index generated without recursive decomposition and clustering;

Fig. 11 shows an example of logic flow for a k-nearest neighbor search using the index generated with recursive decomposition and clustering;

Fig 12 shows an example of data in a 3 dimensional space, and a comparison of the results of a clustering technique based on Euclidean distance and of a clustering technique that adapts to the local structure of the data;

Fig. 13 shows an example of logic flow of a clustering technique that adapts to the local structure of the data;

Fig. 14 shows an example of a complex hyper surface in a 3-dimensional space and two successive approximations generated using a 3-dimensional quad tree generation algorithm; and

Fig. 15 shows an example of logic flow of the determination of the clusters that can contain elements closer than a fixed distance from a given vector, using successive approximations of the geometry of the clusters.

Detailed Description

[0027] Figure 1 depicts an example of a client/server architecture having features of the present invention. As depicted, multiple clients (101) and multiple servers (106) are interconnected by a network (102). The server (106) includes a database management system (DBMS) (104) and direct access storage device (DASD) (105). A query is typically prepared on the client (101) machine and submitted to the server (106) through the network (102). The query typically includes specified data, such as a user provided example or search template, and interacts with a database management system (DBMS) (104) for retrieving or updating a database stored in the DASD (105). An example of a DBMS is that sold by IBM under the trademark DB2.

[0028] According to one aspect of the present invention, queries needing multidimensional, e.g., spatial indexing (including range queries and nearest neighbor queries) will invoke the multidimensional indexing engine (107). The multidimensional indexing engine (107) (described with reference to Figures 8-11) is responsible for retrieving those vectors or records which satisfy the constraints specified by the query based on one or more compact multidimensional indexes (108) and clustering (111) and dimensionality reduction information (112) generated by the index generation logic (110) of the present invention (described with reference to Figures 6 and 7). Most if not all of the compact indexes (108) of the present invention can preferably be stored in the main memory and/or cache of the server (106). Those skilled in the art will appreciate that a database, such as a spatial database, can reside on one or more systems. Those skilled in the art will also appreciate that the multidimensional indexing engine (107) and/or the index generation logic (110) can be combined or incorporated as part of the DBMS (104). The efficient index generation logic (110) and the multidimensional index engine (also called search logic) can be tangibly embodied as software on a computer program product executable on the server (106).

[0029] One example of an application is for stored point-of-sale transactions of a supermarket including geographic coordinates (latitude and longitude) of the store location. Here, the server (106) preferably can also support decision support types of applications to discover knowledge or patterns from the stored data. For example, an online analytical processing (OLAP) engine (103) may be used to intercept queries that are OLAP related, to facilitate their processing. According to the present invention, the OLAP engine, possibly in conjunction with the DBMS, uses the multidimensional index engine (107) to search the index (108) for OLAP related queries. Those skilled in the art will appreciate that the index generation logic (110) of the present invention is applicable to multidimensional data cube representations of a data warehouse. An example of a method and system for generating multidimensional data cube representations of

a data warehouse is described in co-pending United States patent application S/N 08/843,290, filed April 14, 1997, entitled "System and Method for Generating Multi-Representations of a Data Cube," by Castelli et al., which is hereby incorporated by reference in its entirety.

[0030] Multimedia data is another example of data that benefits from spatial indexing. Multimedia data such as audio, video and images can be stored separately from the meta-data used for indexing. One key component of the meta data that can be used for facilitating the indexing and retrieval of the media data are the feature vectors extracted from the raw data. For example, a texture, color histogram and shape can be extracted from regions of the image and be used for constructing indices (108) for retrieval.

[0031] An example of an image mining application is QBIC, which is the integrated search facility in IBM's DB2 Image Extender. QBIC includes an image query engine (server), and a sample client consisting of an HTML graphical user interface and related common gateway interface (CGI) scripts that together form the basis of a complete application. Both the server and the client are extensible so that a user can develop an application-specific image matching function and add it to QBIC. The image search server allows queries of large image databases based on visual image content. It features:

[0032] Querying in visual media. "Show me images like this one", where you define what "like" means in terms of colors, layout, texture, and so on.

[0033] Ranking of images according to similarity to the query image.

[0034] Automatic indexing of images, where numerical descriptors of color and texture are stored. During a search, these properties are used to find similar images.

[0035] The combination of visual queries with text queries or queries on parameters such as a date.

[0036] Similarly, indexes can be generated by first creating a representation of the database to be indexed as a set of vectors, where each vector corresponds to a row in the database and the elements of each vector correspond to the values, for the particular row, contained in the columns for which an index must be generated.

[0037] Creating a representation of the database as a set of vectors is well known in the art. The representation can be created by, but is not limited to: creating for each row of the database an array of length equal to the dimensionality of the index to be generated; and copying to the elements of the array, the values contained in the columns, of the corresponding row, for which the index must be generated.

[0038] Assuming that the i th element of a vector \mathbf{v} is v_i , the vector \mathbf{v} can then be expressed as

$$\mathbf{v} = [v_1 \dots v_N] \quad (1)$$

where 'N' is the number of dimensions of the vector that is used for indexing.

[0039] The client side can usually specify three types of queries, all of which require some form of spatial indexing, as described in the Background:

(1) Exact queries: where a vector is specified and the records or multimedia data that match the vector will be retrieved;

(2) *Range queries*: where the lower and upper limit of each dimension of the vector is specified.

(3) *Nearest neighbor queries*: where the most similar vectors are retrieved based on a similarity measure.

[0040] The most commonly used similarity measure between two vectors, \mathbf{v}_1 and \mathbf{v}_2 , is the Euclidean distance measure, d , defined as

$$d^2 = \sum (v1_i - v2_i)^2 \quad (2)$$

Note that it is not necessary for all of the dimensions i to participate in the computation of either a range or nearest neighbor query. In both cases, a subset of the dimensions can be specified to retrieve the results.

[0041] Figure 2 shows an example of the distribution of the vectors in a multidimensional space. As depicted, a total of three dimensions are required to represent the entire space. However, only two dimensions are required to represent each individual cluster, as cluster 201, 202, and 203 are located on the x-y, y-z, and z-x planes, respectively. Thus, it can be concluded that dimension reduction can be achieved through proper clustering of the data. The same dimensional reduction cannot be achieved by singular value decomposition alone, which can only re-orient the feature space so that the axis in the space coincides with the dominant dimensions (three in this example).

[0042] Eliminating one or more dimensions of a vector is equivalent to projecting the original points into a subspace. Equation (2) shows that only those dimensions where the individual elements in the vector are different, need to be

computed. As a result, the projection of the vector into a subspace does not affect the computation of the distance, providing those elements that are eliminated do not vary in the original space.

[0043] Figure 3 shows an example of a distance computation in the original space and a projected subspace where the projection preserves the relative distance between any two of the three points. As shown, in the original 3-dimensional (3-D) space, a point (301) is more distant from another point (302) than it is from a third point (303). Here, the projections of these points (304, 305 and 306 respectively) into the 2-dimensional (2-D) subspace, preserves the relative distance between points.

[0044] Figure 4 shows an example of projecting three points in a 3-D space into a 2-D space where the projection affects the rank of the relative distance. As shown, the distance between points (401) and (402) in the 3-D space is larger than that between points (402) and (403). In this case however, the distance between the respective projected points (404) and (405) is shorter than the projected distance between (405) and (406). Thus, the relative distance between two points may not be preserved in the projected subspace.

[0045] In the following, a methodology will be derived to estimate the maximum error that can result from projecting vectors into a subspace. The process starts by determining the bound of the maximum error. Denoting the centroid of a cluster as V_c , which is defined as

$$V_c = \frac{1}{N} \sum V_i \quad (3)$$

where N is the total number of vectors in the cluster, which consists of vectors $\{V_1, \dots, V_N\}$. After the cluster is projected into a k dimensional subspace, where without loss of generality the last $(n-k)$ dimensions are eliminated, an error is introduced to the distance between any two vectors in the subspace as compared to the original space. The error term is

$$Error^2 = \sum_{i=k+1}^n (V_{1,i} - V_{2,i})^2 \quad (4)$$

[0046] The following inequality immediately hold:

$$\begin{aligned} Error^2 &\leq \sum_{i=k+1}^n (|V_{1,i}| + |V_{2,i}|)^2 \\ &\leq \sum_{i=k+1}^n (2\max(|V_{1,i}|, |V_{2,i}|))^2 \\ &\leq 4 \sum_{i=k+1}^n \max(|V_{1,i}|, |V_{2,i}|)^2 \end{aligned} \quad (5)$$

[0047] Equation (5) shows that the maximum error incurred by computing the distance in the projected subspace is bounded.

[0048] Figure 5 shows an example of an approximation in computing the distances in accordance with the present invention. The distance between a template point T (501) and a generic point v (506) is given by equation 2. This Euclidean distance is invariant with respect to: rotations of the reference coordinate system; translations of the origin of the coordinate system; reflections of the coordinate axes; and the ordering of the coordinates. Let then, without loss of generality, the point V (506) belong to Cluster 1 in Figure 5. Consider then the reference coordinate system defined by the eigenvectors of the covariance matrix of Cluster 1, and let the origin of the reference coordinate system be Centroid 1. The distance between the template T (501) and a generic point v in cluster 1 (505) can be written as

$$d^2 = \sum_{i=1}^n (T_i - V_i)^2 \quad (6)$$

where the coordinates T_i and V_i are relative to the reference system.

Next, project the elements of Cluster 1 (505) onto Subspace 1, by setting the last $n-k+1$ coordinates to 0. The distance between the template T (501) and the projection v' (507) of v (506) onto Subspace 1 is now

$$\begin{aligned} d_1^2 &= \sum_{i=1}^k (T_i - V_i)^2 + \sum_{i=k+1}^n (T_i)^2 \\ &= d_1'^2 = d_2^2 \end{aligned} \quad (7)$$

[0049] The term d_1 is the Euclidean distance between the projection (504) of the template onto Subspace 1, called Projection 1, and the projection v' (507) of V (506) onto Subspace 1; the term d_2 is the distance between the template T (501) and Projection 1 (504), its projection onto Subspace 1; in other words, d_2 is the distance between the template T (501) and Subspace 1. The approximation introduced can now be bound by substituting equation (7) for equation (6) in the calculation of the distance between the template T (501) and the vector v (506). Elementary geometry teaches that the three points here, T (501), V (506) and V' (507), identify an unique 2-dimensional subspace (a plane). To simplify the discussion, let this plane correspond to the plane (520) shown in Figure 5. Then the distance d defined in equation (6) is equal to the length of the segment joining T (501) and V (506), the distance d' defined in equation (7) is the length of the segment joining T (501) and V' (507). A well known theorem of elementary geometry says that the length of a side of a triangle is longer than the absolute value of the difference of the lengths of the other two sides, and shorter than their sum. This implies that the error incurred in substituting the distance d' defined in equation (7) for the distance d defined in equation (6) is less than or equal to the length of the segment joining V (506) and V' (507); thus, the error squared is bounded by

$$\text{error}^2 \leq \sum_{i=k+1}^n (V_i)^2 \quad (8)$$

[0050] Figure 6 shows an example of a flow chart for generating a hierarchy of reduced dimensionality clusters and low-dimensionality indexes for the clusters at the bottom of the hierarchy. In step 601, a clustering process takes the original data (602) as input; partitions the data into data clusters (603); and generates clustering information (604) on the details of the partition. Each entry in the original data includes a vector attribute as defined in Equation (1). The clustering algorithm can be chosen among, but it is not limited to, any of the clustering or vector quantization algorithms known in the art, as taught, for example, by Leonard Kauffman and Peter J. Rousseeuw in the book "Finding Groups in Data" John Wiley & sons, 1990; or by Yoseph Linde, Andres Buzo and Robert M. Gray in "An Algorithm for Vector Quantizer Design" published in the *IEEE Transactions on Communications*, Vol. COM-28, No. 1, January 1980, P. 84-95. The clustering information (604) produced by the learning stage of a clustering algorithm varies with the algorithm; such information allows the classification stage of the algorithm to produce the cluster number of any new, previously unseen sample, and to produce a representative vector for each cluster. Preferably, the clustering information includes the centroids of the clusters, each associated with a unique label (see e.g., Yoseph Linde, Andres Buzo and Robert M. Gray in "An Algorithm for Vector Quantizer Design", supra). In step 605 a sequencing logic begins which controls the flow of operations so that the following operations are applied to each of the clusters individually and sequentially. Those skilled in the art will appreciate that in the case of multiple computation circuits, the sequencing logic (605) can be replaced by a dispatching logic that initiates simultaneous computations on different computation circuits, each operating on a different data cluster. In step 606, the dimensionality reduction logic (606) receives a data cluster (603) and produces dimensionality reduction information (607) and a reduced dimensionality data cluster (608). In step 609, a termination condition test is performed (discussed below). If the termination condition is not satisfied, steps 601 - 609 can be applied recursively by substituting, in step 611, the reduced dimensionality data cluster (608) for the original data (602) and the process returns to step 601. If in step 610 the termination condition is satisfied, a searchable index (612) is generated for the cluster. In step 613, if the number of clusters already analyzed equals the total number of clusters (603) produced by the clustering algorithm in step 601, the computation terminates. Otherwise, the process returns to step 605. The selection of the number of clusters is usually performed by the user, but there are automatic procedures known in the art; see e.g., Brian Everitt, "Cluster Analysis," Halsted Press, 1973, Chapter 4.2.

[0051] An example of the termination condition test of step 609 can be based on a concept of data volume, $v(x)$,

defined as

$$V(X) = \sum_{i=1}^m n_i$$

where X is a set of records, n_i is the size of the i th record and the sum is over all the elements of X . If the records have the same size S before the dimensionality reduction step 606 and n denotes the number of records in a cluster, then the volume of the cluster is Sn . If S' denotes the size of a record after the dimensionality reduction step 606, then the volume of the cluster after dimensionality reduction is $S'n$. The termination condition can be tested by comparing the volumes Sn and $S'n$ and terminating if $Sn = S'n$.

[0052] In another embodiment, the termination test condition in step 609 is not present and no recursive application of the method is performed.

[0053] Figure 7 shows an example of the dimensionality reduction logic of step 606. As depicted, in step 701, a singular value decomposition logic (701) takes the data cluster (702) as input and produces a transformation matrix (703) and the eigenvalues (704) (or *characteristic values*) of the transformation matrix (703). The columns of the transformation matrix are the eigenvectors (or *characteristic vectors*) of the matrix. Algorithms for singular value decomposition are well known in the art; see e.g., R. A. Horn and C. R. Johnson, "Matrix Analysis" Cambridge University Press (1985). Those skilled in the art will appreciate that the singular value decomposition logic can be replaced by any logic performing a same or equivalent operation. In an alternative embodiment, the singular value decomposition logic can be replaced by a principal component analysis logic, as is also well known in the art.

[0054] In step 705, a sorting logic takes the eigenvalues (704) as input and produces eigenvalues sorted by decreasing magnitude (706). The sorting logic can be, any of the numerous sorting logic well known in the art. In step 707, the selection logic, selects a subset of the ordered eigenvalues (706) containing the largest eigenvalues (708) according to a selection criterion. The selection criterion can be, but it is not limited to, selecting the smallest group of eigenvalues, the sum of which is larger than a user-specified percentage of the trace of the transformation matrix (703), where the trace is the sum of the elements on the diagonal, as is known in the art. In this example, the transformation matrix (703) and the selected eigenvalues (708) comprise dimensionality reduction information (607). Alternatively, the selection of the eigenvalues can be based on a precision vs. recall tradeoff (described below).

[0055] In step 709, the transformation logic takes the data cluster (702) and the transformation matrix (703) as input; and applies a transformation specified by the transformation matrix (703) to the elements of the data cluster (702) and produces a transformed data cluster (710). In step 711, the selected eigenvalues (708) and the transformed data cluster (710) are used to produce the reduced dimensionality data cluster (712). In a preferred embodiment, the dimensionality reduction is accomplished by retaining the smallest number of dimensions such that the set of corresponding eigenvalues account for at least a fixed percentage of the total variance, where for instance the fixed percentage can be taken to be equal to 95%.

[0056] Alternatively, the selection of the eigenvalues can be based on a precision vs. recall tradeoff. To understand precision and recall, it is important to keep in mind that the search operation performed by a method of the current invention can be approximate, (as will be described with reference to Figures 10-11). Let k be the desired number of nearest neighbors to a template in a database of N elements. Here, since the operation is approximate, a user typically requests a number of returned results greater than k . Let n be the number of returned results; of the n results, only c will be correct, in the sense that they are among the k nearest neighbors to the template. The precision is the proportion of the returned results that are correct, and is defined as

$$\text{precision} = c / n,$$

and recall is the proportion of the correct results that are returned by the search, that is,

$$\text{recall} = c / k.$$

[0057] Since precision and recall vary with the choice of the template, their expected values are a better measure of the performance of the system. Thus, both precision and recall are meant as expected values (E) taken over the distribution of the templates, as a function of fixed values of n and k :

$$\text{precision} = E(c)/n,$$

$$\text{recall} = E(c)/k.$$

5

[0058] Note that as the number of returned results n increases, the precision decreases while the recall increases. In general, the trends of precision and recall are not monotonic. Since $E(c)$ depends on n , an efficiency vs. recall curve is often plotted as a parametric function of n . In a preferred embodiment, a requester specifies the desired precision of the search and a lower bound on the allowed recall. Then the dimensionality reduction logic performs the dimensionality reduction based on precision and recall as follows: after ordering the eigenvalues in decreasing order, the dimensionality reduction logic (step 606, Figure 6) removes the dimension corresponding to the smallest eigenvalue, and estimates the resulting precision vs. recall function based on a test set of samples selected at random from the original training set or provided by the user. From the precision vs. recall function, the dimensionality reduction logic derives a maximum value of precision n_{\max} for which the desired recall is attained. Then the dimensionality reduction logic iterates the same procedure by removing the dimension corresponding to the next smallest eigenvalue, and computes the corresponding precision for which the desired recall is attained. The iterative procedure is terminated when the computed precision is below the threshold value specified by the user, and the dimensionality reduction logic retains only the dimensions retained at the iteration immediately preceding the one where the termination condition occurs.

[0059] In another embodiment of the present invention, the requester specifies only a value of desired recall, and the dimensionality reduction logic estimates the cost of increasing the precision to attain the desired recall. This cost has two components: one that decreases with the number of dimensions, since computing distances and searching for nearest neighbors is more efficient in lower dimensionality spaces; and an increasing component due to the fact that the number of retrieved results must grow as the number of retained dimensions is reduced to insure the desired value of recall. Retrieving a larger number n of nearest neighbors is more expensive even when using efficient methods, since the portion of the search space that must be analyzed grows with the number of desired results. Then the dimensionality reduction logic finds by exhaustive search, the number of dimensions to retain that minimizes the cost of the search for the user-specified value of the recall.

[0060] The clustering and singular value decomposition can be applied to the vectors recursively (step 601-611) until a terminating condition (step 609) is reached. One such terminating condition can be that the dimension of each cluster can no longer be reduced as described herein. Optionally, more conventional spatial indexing techniques such as the R-tree can then be applied to each cluster. These techniques are much more efficient for those clusters whose dimension have been minimized. This would thus complete the entire index generation process for a set of high dimensional vectors.

[0061] As will be described with reference to Figures 8-15, the present invention also has features for performing efficient searches using compact representations of multidimensional data. Those skilled in the art will appreciate that the search methods of the present invention are not limited to the specific compact representations of multidimensional data described herein.

[0062] Figure 8 shows an example of a logic flow for an exact search process based on a searchable index (108 or 612) generated according to the present invention. In this example, the index is generated without recursive application of the clustering and singular value decomposition. An exact search is the process of retrieving a record or records that exactly match a search query, such as a search template. As depicted, in step 802 a query including specified data such as a search template (801) is received by the multidimensional index engine (107) (also called cluster search logic). In step 802, the clustering information (604) produced in step 601 of Figure 6, is used to identify the cluster to which the search template belongs. In step 803, the dimensionality reduction information (607) produced in step 606 of Figure 6, is used to project the input template onto the subspace associated with the cluster identified in step 802, and produce a projected template (804). In step 805, an intra-cluster search logic uses the searchable index (612) generated in step 610 of Figure 6 to search for the projected template. Note that the simplest search mechanism within each cluster is to conduct a linear scan (or linear search) if no spatial indexing structure can be utilized. On most occasions, spatial indexing structures such as R-trees can offer better efficiency (as compared to linear scan) when the dimension of the cluster is relatively small (smaller than 10 in most cases).

[0063] Figure 9 shows another example of a flow chart of an exact search process based on a searchable multidimensional index (108 or 612) generated according to the present invention. Here, the index (108 or 612) was generated with a recursive application of the clustering and dimensionality reduction logic. An exact search is the process of retrieving the record or records that exactly match the search template. As depicted, a query including specified data such as a search template (901) is used as input to the cluster search logic, in step 902, which is analogous to the cluster search logic in step 802 of Figure 8. In step 902, clustering information (604) produced in step (601) of Figure

6 is used to identify the cluster to which the search template (901) belongs. In step 903, (analogous to step 803 of Figure 8) the dimensionality reduction information (607), produced in step 606 of Figure 6, is used to project the input template onto the subspace associated with the cluster identified in step 902, and produce a projected template (904). In step 905, it is determined whether the current cluster is terminal, that is, if no further recursive clustering and singular value decomposition steps were applied to this cluster during the multidimensional index construction process. If the cluster is not terminal, in step 907 the search template (901) is replaced by the projected template (904), and the process returns to step 902. In step 906, if the cluster is terminal, the intra-cluster search logic uses the searchable index to search for the projected template. As noted, the simplest intra-cluster search mechanism is to conduct a linear scan (or linear search), if no spatial indexing structure can be utilized. On most occasions, spatial indexing structures such as R-trees can offer better efficiency (as compared to linear scan) when the dimension of the cluster is relatively small (smaller than 10 in most cases).

[0064] The present invention has also features for assessing if other clusters can contain elements that are closer to the specific data than the farthest of the k most similar elements retrieved. As is known in the art, clustering information can be used to reconstruct the boundaries of the partitions, and these boundaries can be used to determine if a cluster can contain one of the k nearest neighbors. Those skilled in the art will appreciate that the cluster boundaries are a simple approximation to the structure of the cluster itself. In other words, from the mathematical form of the boundary it is not possible to tell whether there are elements of clusters near any given position on the boundary. As an example, consider a case where the database contains two spherical clusters of data which are relatively distant from each other. A reasonable boundary for this case would be a hyperplane, perpendicular to the line joining the centroids of the clusters, and equidistant from the centroids. Since the clusters are relatively distant however, there is no data point near the boundary. In other cases, the boundary could be very close to large number of elements of both clusters.

[0065] As will be described with reference to Figures 14 and 15, the present invention has features computing and storing, in addition to the clustering information, a hierarchy of approximations to the actual geometric structure of each cluster; and using the approximation hierarchy to identify clusters which can contain elements closer than a fixed distance from a given vector.

[0066] Figure 10 shows an example of a flow chart of a k-nearest neighbor search process based on an index (612) generated according to the present invention. In this example, the index is generated without recursive application of the clustering and singular value decomposition. The k-nearest neighbor search returns the k closest entries in the database which match the query. The number (nr) of desired matches, k (1000) is used in step 1002 to initialize the k-nearest neighbor set (1009) so that it contains at most k elements and so that it is empty before the next step begins. In step 1003, the cluster search logic receives the query, for example a the search template (1001) and determines which cluster the search template belongs to, using the clustering information (604) produced in step 601 of Figure 6. In step 1004, template is then projected onto the subspace associated with the cluster it belongs to, using the dimensionality reduction information (607). The projection step 1004 produces a projected template (1006) and dimensionality reduction information (1005), that includes an orthogonal complement of the projected template (1006) (defined as the vector difference of the search template (1001) and the projected template (1006)), and the Euclidean length of the orthogonal complement. The dimensionality reduction information (1005) and the projected template (1006) can be used by the intra-cluster search logic, in step 1007, which updates the k-nearest neighbor set (1009), using the multidimensional index. Examples of the intra-cluster search logic adaptable according to the present invention include any of the nearest neighbor search methods known in the art; see e.g., "Nearest Neighbor Pattern Classification Techniques" Belur V. Desai, (editor), IEEE Computer Society (1991). An example of an intra-cluster search logic (step 1007) having features of the present invention includes the steps of: first, the squared distance between the projected template (1006) and the members of the cluster in the vector space with reduced dimensionality is computed; the result is added to the squared distance between the search template (1001) and the subspace of the cluster; the final result is defined as the "sum" of the squared lengths of the orthogonal complements, (computed in step (1004) which is part of the dimensionality reduction information (1005):

$$\delta^2(\text{template}, \text{element}) = D^2(\text{projected_template}, \text{element}) + \sum \|\text{orthoqonal_complement}\|^2$$

[0067] If the k-nearest neighbor set (1009) is empty at the beginning of step 1007, then the intra-cluster search fills the k-nearest neighbor set with the k elements of the cluster that are closest to the projected template (1006) if the number of elements in the cluster is greater than k, or with all the elements of the cluster otherwise. Each element of the k-nearest neighbor set is associated with a corresponding mismatch index δ^2 .

[0068] If the k-nearest neighbor set (1009), is not empty at the beginning of step 1007, then the intra-cluster search logic, in step 1007 updates the k-nearest neighbor set when an element is found whose mismatch index δ^2 is smaller than the largest of the indexes currently associated with elements in the k-nearest neighbor set (1009). The k-nearest neighbor set can be updated by removing the element with largest mismatch index δ^2 from the k-nearest neighbor set

(1009) and substituting the newly found element for it.

[0069] If the k -nearest neighbor set (1009) contains less than k elements, the missing elements are considered as elements at infinite distance. In step 1008, it is determined whether there is another candidate cluster that can contain nearest neighbors. This step takes the clustering information (604) as input, from which it can determine the cluster boundaries. If the boundaries of a cluster to which the search template (1001) does not belong are closer than the farthest element of the k -nearest neighbor set (1009), then the cluster is a candidate. If no candidate exists, the process terminates and the content of the k -nearest neighbor set (1009) is returned as a result. Otherwise, the process returns to step 1004, where the current cluster is now the candidate cluster identified in step 1008.

[0070] Figure 11 shows an example of a flow chart of a k -nearest neighbor search process based on a search index (612) generated according to the present invention. Here, the index is generated with a recursive application of the clustering and singular value decomposition. The k -nearest neighbor search returns the closest k entries in the database, to specified data in the form of a search template. As depicted, in step 1102, the k -nearest neighbor set is initialized to empty and the number of desired matches, k (1100) is used to initialize the k -nearest neighbor set (1111) so that it can contain at most k elements. In step 1103, the cluster search logic takes the search template (1101) as input and associates the search template to a corresponding cluster, using the clustering information (604) produced in step 601 of Figure 6. In step 1104, the template (1101) is projected onto a subspace associated with the cluster identified in step 1103, based on the dimensionality reduction information (607) produced in step 606 of Figure 6. In step 1104, a projected template (1106) and dimensionality reduction information (1105) are produced. Preferably, the dimensionality reduction information (1105) includes the orthogonal complement of the projected template (1106) (defined as the vector difference of the search template (1101) and the projected template (1106)), and the Euclidean length of the orthogonal complement. In step 1107, it is determined if the current cluster is terminal, that is, if no further recursive steps of clustering and singular value decomposition were applied to the current cluster during the construction of the index. If the current cluster is not terminal, in step 1108 the projected template (1106) is substituted for the search template (1101) and the process returns to step 1103. Otherwise, the dimensionality reduction information (1105) and the projected template (1106) are used by the intra-cluster search logic in step 1109, to update the k -nearest neighbor set (1111), based on the searchable index (612). Examples of intra-cluster search logic which are adaptable according to the present invention include any of the nearest neighbor search methods known in the art; see e.g., "Nearest Neighbor Pattern Classification Techniques" IEEE Computer Society, Belur v. Desathary (editor), 1991. One example of an intra-cluster search logic (step 1007) logic having features of the present invention includes the steps of: first, the squared distance between the projected template (1006) and the members of the cluster in the vector space with reduced dimensionality is computed; the result is added to the squared distance between the search template (1001) and the subspace of the cluster; the result is defined as the "sum" of the squared lengths of the orthogonal complements, (computed in step (1004) which is part of the dimensionality reduction information (1005));

$$\delta^2(template, element) = D^2(projected_template, element) + \sum \|orthogonal_complement\|^2$$

[0071] If the k -nearest neighbor set (1111) is empty at the beginning of step 1109, then the intra-cluster search logic fills the k -nearest neighbor set (1111) with either: the k elements of the cluster that are closest to the projected template (1106) if the number of elements in the cluster is greater than k ; or with all the elements of the cluster if the number of elements in the cluster is equal or less than k . Each element of the k -nearest neighbor set (1111) is preferably associated with its corresponding mismatch index δ^2 .

[0072] If the k -nearest neighbor set (1111) is not empty at the beginning of step 1109, then the intra-cluster search logic updates the k -nearest neighbor set when an element is found whose mismatch index is smaller than the largest of the indexes currently associated with the elements in the k -nearest neighbor set (1111). The update can comprise removing the element with the largest mismatch index δ^2 from the k -nearest neighbor set (1111) and substituting the newly found element for it.

[0073] If the k -nearest neighbor set (1111) contains less than k elements, the missing elements are considered as elements at infinite distance. In step 1110, it is determined whether the current level of the hierarchy is the top level (before the first clustering step is applied). If the current level is the top level, then the ends and the content of the k -nearest neighbor set (1111) is returned as a result. If the current level is not the top level, then in step 1114, a search is made for a candidate cluster at the current level, that is, a cluster that can contain some of the k ' nearest neighbors. The search is performed using the dimensionality reduction information (1105) and the clustering information (604). In step 1114, the clustering information (604) is used to determine the cluster boundaries. If the boundaries of a cluster to which the search template (1001) does not belong are closer than the farthest element of the k -nearest neighbor set (1111), then the cluster is a candidate. If no candidate exists, then in step 1113, the current level is set to the previous level of the hierarchy and the dimensionality reduction information is updated, and the process returns to step 1110. If a candidate exists, then in step 1115, the template is projected onto the candidate cluster, thus updating the projected

template (1106); and the dimensionality reduction information is updated. The process then returns to step 1107.

[0074] Figures 12 (a)-12 (c) compares the results of clustering techniques based on similarity alone (for instance, based on the minimization of the Euclidean distance between the element of each cluster and the corresponding centroids as taught by Linde, Buzo and Gray in "An Algorithm for vector Quantizer Design" supra), with clustering using an algorithm that adapts to the local structure of the data. Figure 12 (a) shows a coordinate reference system (1201) and a set of vectors (1202). If a clustering technique based on minimizing the Euclidean distance between the elements of each cluster and the corresponding centroid is used, a possible result is shown in Figure 12 (b): the set of vectors (1202) is partitioned into three clusters, cluster 1 (1205), cluster 2 (1206) and cluster 3 (1207) by the hyper planes (1203) and (1204). The resulting clusters contain vectors that are similar to each other, but do not capture the structure of the data, and would result in sub-optimal dimensionality reduction. Figure 12 (c) shows the results of clustering using an algorithm that adapts to the local structure of the data. This results in three clusters, cluster 1 (1208), cluster 2 (1209) and cluster 3 (1210), that better capture the local structure of the data and are more amenable to independent dimensionality reduction.

[0075] Figure 13 shows an example of a clustering algorithm that adapts to the local structure of the data. In step 1302, the set of vectors (1301) to be clustered and the desired number of clusters are used to select initial values of the centroid (1303). In a preferred embodiment, one element of the set of vectors is selected at random for each of the desired number *NC* of clusters, using any of the known sampling without replacing techniques. In step 1304, a first set of clusters is generated, for example using any of the known methods based on the Euclidean distance. As a result, the samples are divided into *NC* clusters (1305). In step 1306, the centroids (1307) of each of the *NC* clusters are computed, for instance as an average of the vectors in the cluster. In step 1308, the eigenvalues and eigenvectors (1309) of the clusters (1305) can be computed using the singular value decomposition logic (Figure 7, step 701). In step 1310, the centroid information (1307), the eigenvectors and eigenvalues (1309) are used to produce a different distance metric for each cluster. An example of the distance metric for a particular cluster is the weighted Euclidean distance in the rotated space defined by the eigenvectors, with weights equal to the square root of the eigenvalues.

[0076] The loop formed by logic steps 1312, 1313 and 1314 generates new clusters. In step 1312, a control logic iterates step 1313 and 1314 over all the vectors in the vector set (1301). In step 1313, the distance between the selected example vector and each of the centroids of the clusters is computed using the distance metrics (1311). In step 1314, the vector is assigned to the nearest cluster, thus updating the clusters (1305). In step 1315, if a termination condition is reached, the process ends; otherwise the process continues at step 1306. In a preferred embodiment, the computation is terminated if no change in the composition of the clusters occurs between two subsequent iterations.

[0077] Figure 14 shows an example of a complex surface (1401) in a 3-dimensional space and two successive approximations (1402, 1403) based on a 3-dimensional quad tree, as taught in the art by Samet, H. in "Region Representation Quadtree from Boundary Codes" Comm. ACM 23,3, pp. 163-170 (March 1980). The first approximation (1402) is a minimal bounding box. The second approximation (1403) is the second step of a quad tree generation, where the bounding box has been divided into 8 hyper rectangles by splitting the bounding box at the midpoint of each dimension, and retaining only those hyper rectangles that intersect the surface.

[0078] In a preferred embodiment, the hierarchy of approximations is generated as a *k*-dimensional quad tree. An example of a method having features of the present invention for generating the hierarchy of approximations includes the steps of: generating the cluster boundaries, which correspond to a zeroth-order approximation to the geometry of the clusters; approximating the convex hull of each of the clusters by means of a minimum bounding box, thus generating a first-order approximation to the geometry of each cluster; partitioning the bounding box into 2^k hyper rectangles, by cutting it at the half point of each dimension; retaining only those hyper rectangles that contain points, thus generating a second-order approximation to the geometry of the cluster; and repeating the last two steps for each of the retained hyper rectangles to generate successively the third-, fourth-, ..., *n*-th approximation to the geometry of the cluster.

[0079] Figure 15 shows an example of logic flow to identify clusters that can contain elements closer than a prescribed distance from a given data point, using the hierarchy of successive approximations to the geometry of the clusters. In one embodiment, the geometry of a cluster is a convex hull of the cluster. In another embodiment, the geometry of a cluster is a connected elastic surface that encloses all the points. This logic can be used in searching for candidate clusters, for example in step 1008 of Figure 10. Referring now to Figure 15, in step 1502, the original set of clusters with their hierarchy of geometric approximations (1501) are input to the process. In step 1502, the candidate set (1505) is initialized to the original set (1501). In step 1506, another initialization step is performed by setting the current approximations to the geometry to zero-th order approximations. In a preferred embodiment, the zero-th order geometric approximations of the clusters are given by the decision regions of the clustering algorithm used to generate the clusters. In step 1507, the distances between the current approximations of the geometry of the cluster and the data point (1503) are computed. All clusters more distant than the candidate set (1505) are discarded and a retained set (1508) of clusters is produced. In step 1509, it is determined if there are better approximations in the hierarchy. If no better approximations exist, the computation is terminated by setting the result set (1512) to be equal to the currently retained set (1508). Otherwise, in step 1510, the candidate set (1505) is set to the currently retained set (1508), the current geometric

approximation is set to the immediately better approximation in the hierarchy, and the process return to step 1507.

Claims

1. A computerized method of representing multidimensional data, comprising the steps of:
 - a) partitioning the multidimensional data into one or more clusters;
 - b) generating and storing clustering information for said one or more clusters;
 - c) generating one or more reduced dimensionality clusters and dimensionality reduction information for said one or more clusters; and
 - d) storing the dimensionality reduction information.
2. The method of claim 1, further comprising the steps of:
 - generating and storing a reduced dimensionality index for said one or more reduced dimensionality clusters.
3. The method of claim 1 wherein the data is stored in one of a spatial database or a multimedia database which includes a plurality of data records each containing a plurality of fields, further comprising the steps of:
 - creating a representation of the database to be indexed as a set of vectors, where each vector corresponds to a row in the database and the elements of each vector correspond to the values, for a particular row, contained in the columns for which the searchable index will be generated; and
 - said partitioning comprising partitioning the vectors into said one or more clusters.
4. The method of claim 2, further comprising the step of storing the index entirely in the main memory of a computer.
5. The method of claim 2, wherein said generating a reduced dimensionality cluster comprises a singular value decomposition further comprising the steps of:
 - producing a transformation matrix and eigenvalues of the transformation matrix for said each cluster; and
 - selecting a subset of the eigenvalues including the largest eigenvalues; wherein the dimensionality reduction information includes the transformation matrix and the subset of the eigenvalues.
6. The method of claim 5, for searching for k most similar records to specified data using the reduced dimensionality index, comprising the steps of:
 - associating specified data to said one or more clusters, based on stored clustering information;
 - projecting the specified data onto a subspace for an associated cluster based on stored dimensionality reduction information for the associated cluster;
 - generating dimensionality reduction information including an orthogonal complement for projected specified data, in response to said projecting;
 - searching, via the index, for the associated cluster having k most similar records to the projected specified data;
 - determining if any other associated cluster can include any of k most similar records to the projected specified data; and
 - repeating said searching on said any cluster that can include any of the k most similar records to the specified data.
7. The method of claim 6, wherein the specified data includes a search template, further comprising the steps of:
 - said projecting step including projecting the template, using the dimensionality reduction information, onto a subspace associated with the cluster to which it belongs;
 - generating template dimensionality reduction information for a projected template; wherein said searching, via the index is based on the projected template and the template dimensionality reduction information; and
 - updating a k-nearest neighbor set of the k most similar records to the search template.
8. The method of claim 5, wherein said selecting a subset of the eigenvalues is a function of a precision and a recall of returned results.

9. The method of claim 2, for searching for k records most similar to specified data, the method for searching comprising the steps of:

identifying a cluster to which specified data belongs, based on the clustering information;
 reducing the dimensionality of the specified data based on the dimensionality reduction information for an identified cluster;
 generating dimensionality reduction information for reduced dimensionality specified data, in response to said reducing;
 searching the multidimensional index, using the dimensionality reduction information, for a reduced-dimensionality version of the cluster to which the specified data belongs;
 retrieving via the multidimensional index, the k most similar records in the cluster;
 identifying other candidate clusters which can contain records closer to the specified data than the farthest of the k most similar records retrieved;
 searching a closest other candidate cluster to the specified data, in response to said determining step; and
 repeating said identifying and searching steps for all said other candidate clusters.

10. The method of claim 6 or 9, further comprising the step of:

computing the distances (D) between k nearest neighbors in the version of the cluster and the projected specified data as a function of a mismatch index δ^2 wherein,

$$\delta^2(template, element) = D^2(projected_template, element) + \sum \|orthoqonal_complement\|^2.$$

11. The method of claim 1, wherein the clustering information includes information on a centroid of said one or more clusters, further comprising the step of associating the centroid with a unique label.

12. The method of claim 1, wherein the dimensionality of the data is > 8.

13. The method of claim 1, for performing an exact search, comprising the steps of:

associating specified data to one of the clusters based on stored clustering information;
 reducing the dimensionality of the specified data based on stored dimensionality reduction information for a reduced dimensionality version of a cluster, in response to said associating step; and
 searching, based on reduced dimensionality specified data, for the reduced dimensionality version of the cluster matching the specified data.

14. The method of claim 13, wherein said searching further comprises the step of conducting a linear scan to match the specified data.

15. The method of claim 1, further comprising the steps of:

creating a hierarchy of reduced dimensionality clusters by recursively applying said steps a) through d); and
 generating and storing one or more low-dimensional indexes for clusters at a lowest level of said hierarchy.

16. The method of claim 15, for performing an exact search, comprising the steps of:

recursively applying the steps of:

finding a cluster to which specified data belongs, using stored clustering information; and
 reducing the dimensionality of the specified data using stored dimensionality reduction information, until a corresponding lowest level of the hierarchy of reduced dimensionality clusters has been reached; and
 searching, using the low dimensional indexes, for a reduced dimensionality version of the cluster matching the specified data.

17. The method of claim 15, for performing a similarity search, further comprising the steps of:

recursively applying the steps of:

finding the cluster to which specified data belongs, using stored clustering information; and
 reducing the dimensionality of the specified data to correspond to the lowest level of the hierarchy of
 reduced dimensionality clusters, using stored dimensionality reduction information;
 searching for candidate terminal clusters that can contain one or more of k nearest neighbors of the spec-
 5 ified data at each level of the hierarchy of reduced dimensionality clusters starting from a terminal cluster
 at a lowest level of said hierarchy to which the specified data belongs; and
 for each candidate terminal cluster, performing an intra-cluster search for the k nearest neighbors to the
 specified data.

10 **18.** The method of claim 15, for performing a similarity search, further comprising the steps of:

reducing the dimensionality of the specified data;
 recursively applying the steps of: finding the cluster to which reduced dimensionality specified data belongs,
 using stored clustering information; and reducing the dimensionality of the reduced dimensionality specified
 15 data to correspond to a lowest level of a hierarchy of reduced dimensionality clusters, using stored dimen-
 sionality reduction information;
 searching for candidate terminal clusters that can contain one or more of k nearest neighbors of the reduced
 dimensionality specified data at each level of the hierarchy of reduced dimensionality clusters starting from a
 terminal cluster at a lowest level of said hierarchy to which the specified data belongs; and
 20 for each candidate terminal cluster, performing an intra-cluster search for the k nearest neighbors to the re-
 duced dimensionality specified data.

19. The method of claim 1, wherein the data is stored in a database, further comprising the steps of:

25 reducing a dimensionality of the database and generating dimensionality reduction information associated
 with the database; and
 storing the dimensionality reduction information associated with the database;

wherein said partitioning step is responsive to said reducing step.

30 **20.** The method of claim 19, for performing an exact search, comprising the steps of:

reducing the dimensionality of specified data, based on the dimensionality reduction information for the data-
 base;
 35 associating reduced dimensionality specified data to one of the clusters, based on the clustering information,
 in response to said reducing;
 reducing a dimensionality of the specified data to that of a reduced dimensionality cluster defined by an as-
 sociated cluster, based on dimensionality reduction information for the associated cluster; and
 searching for a matching reduced dimensionality cluster, based on a reduced dimensionality version the spec-
 40 ified data.

21. The method of claim 19, for performing a similarity search, comprising the steps of:

45 reducing the dimensionality of specified data using the dimensionality reduction information associated with
 the database;
 finding the cluster to which reduced dimensionality specified data belongs, based on the clustering information;
 reducing the dimensionality of the reduced dimensionality specified data based on the dimensionality reduction
 information for an identified cluster;
 searching for a reduced-dimensionality version of the cluster to which the further reduced dimensionality spec-
 50 ified data belongs;
 retrieving via the multidimensional index, k records most similar to the further reduced dimensionality specified
 data in the cluster;
 assessing whether other clusters can contain records closer to the specified data than the farthest of the k
 records retrieved;
 55 searching a closest other cluster to the specified data, in response to said assessing step; and
 repeating said assessing and searching for all said other clusters.

22. The method of claim 19, wherein the data is stored in a database, further comprising the step of: generating and

storing one or more reduced dimensionality searchable indexes for said one or more reduced dimensionality clusters.

23. The method of claim 19, for performing an exact search, comprising the steps of:

associating specified data to one of the clusters based on stored clustering information;
decomposing the specified data into a reduced dimensionality cluster defined by an associated cluster and stored dimensionality reduction information for the associated cluster, in response to said associating; and
searching said indexes for a matching reduced dimensionality cluster based on decomposed specified data.

24. The method of claim 23, wherein the query includes a search template, further comprising the steps of:

said associating comprising identifying the search template with a cluster based on the stored clustering information;
said decomposing comprising projecting the search template onto a subspace for an identified cluster based on the stored dimensionality reduction information; and
said searching comprising performing an intra-cluster search for a projected template.

25. The method of claim 1, further comprising the steps of:

(a) generating cluster boundaries, which correspond to a zeroth-order approximation to a geometry of said clusters;
(b) approximating the geometry of each of the clusters by means of a minimum bounding box and generating therefrom a first-order approximation to the geometry of each cluster;
(c) partitioning the bounding box into $2k$ hyper rectangles, wherein said partitioning is at a midpoint of each dimension;
(d) retaining only those hyper rectangles that contain data points and generating therefrom a second-order approximation to the geometry of the cluster; and
(e) repeating said steps (c) and (d) for each retained hyper rectangle to generate successively the third-, fourth-, ..., n -th approximations to the geometry of the cluster.

26. The method of claim 25, for searching a hierarchy of approximations to the geometric structure of each cluster, further comprising the steps of:

reducing the dimensionality of the specified data using the dimensionality reduction information associated with the database;
finding the cluster to which reduced dimensionality specified data belongs, based on the clustering information;
reducing the dimensionality of the reduced dimensionality specified data based on the dimensionality reduction information for an identified cluster;
searching for a reduced-dimensionality version of the cluster to which the further reduced dimensionality specified data belongs;
retrieving via the multidimensional index, k records most similar to the further reduced dimensionality specified data in the cluster;
assessing whether one or more other clusters can contain records closer to the specified data than the farthest of the k records retrieved;
retaining an other cluster only if it can contain any of k -nearest neighbors of specified data based on boundaries of the cluster;
iteratively determining if a retained cluster could contain any of the k -nearest neighbors, based on increasingly finer approximations to the geometry, and retaining the retained cluster only if the cluster is accepted at the finest level of the hierarchy of successive approximations; and
identifying a retained cluster as a candidate cluster containing one or more of the k -nearest neighbors of the data, in response to said step of iteratively determining.

27. A program storage device readable by a machine which includes one or more reduced dimensionality indexes to multidimensional data, the program storage device tangibly embodying a program of instructions executable by the machine to perform method steps for representing multidimensional data as claimed in claim 1.

28. A computer program product comprising:

a computer usable medium having computer readable program code means embodied therein for representing multidimensional data, the computer readable program code means in said computer program product comprising:

computer readable program code clustering means for causing a computer to effect partitioning the multidimensional data into one or more clusters;
 computer readable program code means, coupled to said clustering means, for causing a computer to effect generating and storing clustering information for said one or more clusters;
 computer readable program code dimensionality reduction means, coupled to said clustering means, for causing a computer to effect, generating one or more reduced dimensionality clusters and dimensionality reduction information for said one or more clusters; and
 computer readable program code means, coupled to said dimensionality reduction means, for causing a computer to effect storing the dimensionality reduction information.

Patentansprüche

1. Ein computerisiertes Verfahren zur Darstellung von mehrdimensionalen Daten, das Schritte enthält, um

- a) die mehrdimensionalen Daten in einem oder mehreren Cluster zu partitionieren;
- b) Gruppierungsinformationen für einen oder mehrere Cluster zu erstellen und zu speichern;
- c) einen oder mehrere dimensionsverringerte Cluster und Dimensionsverringeringsinformationen für einen oder mehrere Cluster zu erstellen; und
- d) die Dimensionsverringeringsinformationen zu speichern.

2. Das Verfahren nach Anspruch 1, das außerdem Schritte enthält, um

einen dimensionsverringerten Index für einen oder mehrere dimensionsverringerte Cluster zu erstellen und zu speichern.

3. Das Verfahren nach Anspruch 1, wobei die Daten in einer räumlichen Datenbank oder einer Multimedia-Datenbank gespeichert werden, die eine Vielzahl von Datenaufzeichnungen enthält, von denen jede eine Vielzahl von Feldern enthält, wobei das Verfahren außerdem Schritte enthält, um

eine Darstellung von der zu indizierenden Datenbank als einen Satz Vektoren zu erstellen, wobei jeder Vektor einer Zeile in der Datenbank entspricht, und die Elemente von jedem Vektor den Werten für die bestimmte Zeile entsprechen, die in den Spalten enthalten ist, für die der suchbare Index zu erstellen ist; und

die Partitionierung das Partitionieren der Vektoren in einem oder mehreren Cluster enthält.

4. Das Verfahren nach Anspruch 2, das außerdem Schritte enthält, um den Index komplett im Hauptspeicher eines Computers zu speichern.

5. Das Verfahren nach Anspruch 2, wobei die Erstellung eines dimensionsverringerten Clusters eine singularische Wertzerlegung enthält, und das Verfahren außerdem Schritte enthält, um

eine Transformationsmatrix und Eigenwerte von der Transformationsmatrix für jeden Cluster zu erstellen; und

um eine Untermenge mit Eigenwerten auszuwählen, die die größten Eigenwerte enthält; wobei die Dimensionsverringeringsinformation die Transformationsmatrix und die Untermenge der Eigenwerte enthält.

6. Das Verfahren nach Anspruch 5, um in den spezifizierten Daten mittels des dimensionsverringerten Indizes nach Aufzeichnungen zu suchen, die k am ähnlichsten sind, wobei das Verfahren Schritte enthält, um

die spezifizierten Daten einem oder mehreren Cluster auf der Basis der gespeicherten Gruppierungsinformationen zuzuweisen;

die spezifizierten Daten für einen zugehörigen Cluster auf der Basis der für den zugehörigen Cluster gespeicherten Dimensionsverringeringinformationen in einen Unterraum zu projizieren;

Dimensionsverringeringinformationen zu erstellen, die ein orthogonales Komplement für die projizierten spezifizierten Daten als Reaktion auf die Projektion enthält;

über den Index nach dem zugehörigen Cluster, der in den projizierten spezifizierten Daten die Aufzeichnungen enthält, die k am ähnlichsten sind, zu suchen;

zu bestimmen, ob ein anderer zugehöriger Cluster in den projizierten spezifizierten Daten eine der Aufzeichnungen, die k am ähnlichsten sind, enthalten kann; und um

die Suche in jedem Cluster zu wiederholen, der in den spezifizierten Daten eine der Aufzeichnungen enthalten kann, die k am ähnlichsten sind.

7. Das Verfahren nach Anspruch 6, wobei die spezifizierten Daten eine Suchschablone enthalten, und das Verfahren außerdem Schritte enthält, um

den Projektionsschritt, der das Projizieren der Schablone enthält, mittels der Dimensionsverringeringinformationen in einem Unterraum auszuführen, der mit dem Cluster verbunden ist, zu dem er gehört;

die Schablonen-Dimensionsverringeringinformation für eine projizierte Schablone zu erstellen; wobei die Suche über den Index auf der projizierten Schablone und der Schablonen-Dimensionsverringeringinformation basiert; und um

einen k-nearest-neighbor-Satz von den Aufzeichnungen, die k am ähnlichsten sind in der Suchschablone zu aktualisieren.

8. Das Verfahren nach Anspruch 5, wobei die Auswahl einer Untermenge von Eigenwerten abhängig von einer Genauigkeit und einem Recall der zurückgesendeten Ergebnisse ist.

9. Das Verfahren nach Anspruch 2, um in den spezifizierten Daten nach Aufzeichnungen, die k am ähnlichsten sind, zu suchen, wobei das Verfahren zum Suchen Schritte enthält, um

einen Cluster, zu dem die spezifizierten Daten gehören, anhand der Gruppierungsinformation zu identifizieren;

die Dimensionalität der spezifizierten Daten anhand der Dimensionsverringeringinformation für einen identifizierten Cluster zu verringern;

die Dimensionsverringeringinformation für die dimensionsverringerten spezifizierten Daten als Reaktion auf die Verringerung zu erstellen;

den mehrdimensionalen Index mittels der Dimensionsverringeringinformation nach einer dimensionsverringerte Version des Clusters zu durchsuchen, zu dem die angegebenen Daten gehören;

über den mehrdimensionalen Index die Aufzeichnungen, die k am ähnlichsten sind, im Cluster wiederzufinden;

weitere Kandidaten-Cluster zu identifizieren, die Aufzeichnungen enthalten können, die näher an den spezifizierten Daten sind als die weitesten der wiedergefundenen Aufzeichnungen, die k am ähnlichsten sind;

als Reaktion auf den Bestimmungsschritt nach einem weiteren Kandidaten-Cluster zu suchen, der den spezifizierten Daten näher ist; und um

die Schritte zum Identifizieren und Suchen für alle anderen Kandidaten-Cluster zu wiederholen;

10. Das Verfahren nach Anspruch 6 oder 9, das außerdem Schritte enthält, um

Entfernungen (D) zwischen den k-nearest-neighbors in der Version des Clusters und den projizierten spezifizierten Daten als eine Funktion eines Ungleichheitsindizes Δ^2 zu berechnen, wobei

$$\ddot{a}^2(\text{Schablone}, \text{Element}) = D^2(\text{projected_template}, \text{element}) + \|\text{orthogonal_complement}\|^2$$

ist.

11. Das Verfahren nach Anspruch 1, wobei die Gruppierungsinformation Informationen zu einem Schwerpunkt von einem oder mehreren Cluster enthält, wobei das Verfahren außerdem den Schritt enthält, den Schwerpunkt einem einzigartigen Label zuzuweisen.

12. Das Verfahren nach Anspruch 1, wobei die Dimensionalität der Daten > 8 ist.

13. Das Verfahren nach Anspruch 1, um eine genaue Suche durchzuführen, wobei das Verfahren Schritte enthält, um die spezifizierten Daten basierend auf den gespeicherten Gruppierungsinformationen einem der Cluster zuzuordnen;

die Dimensionalität der spezifizierten Daten basierend auf den für eine dimensionsverringerte Version eines Clusters gespeicherten Dimensionsverringierungsinformationen als Reaktion auf den Zuordnungsschritt zu verringern; und um

basierend auf den dimensionsverringerten, spezifizierten Daten nach der dimensionsverringerten Version des Clusters zu suchen, der zu den spezifizierten Daten paßt.

14. Das Verfahren nach Anspruch 13, wobei das Suchen außerdem den Schritt enthält, eine lineare Abtastung durchzuführen, um die spezifizierten Daten abzugleichen.

15. Das Verfahren nach Anspruch 1, das außerdem Schritte enthält, um

eine Hierarchie mit dimensionsverringerten Cluster zu erstellen, indem die Schritte a) bis d) rekursiv angelegt werden; und

einen oder mehrere dimensionsverringerte Indizes für Cluster in der untersten Ebene der Hierarchie zu erstellen und zu speichern.

16. Das Verfahren nach Anspruch 15, um eine genaue Suche durchzuführen, das Schritte enthält, um

die folgenden Schritte rekursiv anzuwenden:

mittels der gespeicherten Gruppierungsinformationen einen Cluster zu finden, zu dem die angegebenen Daten gehören; und

die Dimensionalität der spezifizierten Daten mittels der gespeicherten Dimensionsverringierungsinformationen zu verringern, bis eine entsprechende unterste Ebene in der Hierarchie der dimensionsverringerten Cluster erreicht wurde; und

mittels der dimensionsverringerten Indizes nach einer dimensionsverringerten Version der Cluster zu suchen, die zu den spezifizierten Daten passen.

17. Das Verfahren nach Anspruch 15, um eine Ähnlichkeitssuche durchzuführen, das außerdem Schritte enthält, um

die folgenden Schritte rekursiv anzulegen:

mittels der gespeicherten Gruppierungsinformationen einen Cluster zu finden, zu dem die angegebenen Daten gehören; und

mittels der gespeicherten Dimensionsverringierungsinformation die Dimensionalität der spezifizierten Daten zu verringern, damit sie der untersten Ebene in der Hierarchie der dimensionsverringerten Cluster entsprechen;

nach den abgeschlossenen Kandidaten-Cluster zu suchen, die einen oder mehrere der k-nearest-neighbors in den Daten enthalten können, die in jeder Ebene in der Hierarchie mit dimensionsverringerten Cluster spezifiziert sind, wobei bei einem abgeschlossenen Cluster in der untersten Ebene in dieser Hierarchie, zu der die spezifizierten Daten gehören, begonnen wird; und

für jeden abgeschlossenen Kandidaten-Cluster eine interne Cluster-Suche nach den in den spezifizierten Daten enthaltenen k-nearest-neighbors zu suchen.

18. Das Verfahren nach Anspruch 15, um eine Ähnlichkeitssuche durchzuführen, das außerdem Schritte enthält, um

die Dimensionalität der spezifizierten Daten zu reduzieren;

die folgenden Schritte rekursiv anzuwenden: mittels der gespeicherten Gruppierungsinformation den Cluster zu finden, zu dem die dimensionsverringerten spezifizierten Daten gehören; und mittels der gespeicherten Dimensionsverringierungsinformationen die Dimensionalität der dimensionsverringerten Daten zu verringern, damit sie der untersten Ebene in einer Hierarchie mit dimensionsverringerten Cluster entsprechen;

nach den abgeschlossenen Kandidaten-Cluster zu suchen, die einen oder mehrere der k-nearest-neighbors in den dimensionsverringerten Daten enthalten können, die in jeder Ebene in der Hierarchie mit dimensionsverringerten Cluster spezifiziert sind, wobei bei einem abgeschlossenen Cluster in der untersten Ebene in dieser Hierarchie, zu der die spezifizierten Daten gehören, begonnen wird; und

für jeden abgeschlossenen Kandidaten-Cluster eine interne Cluster-Suche nach den in den dimensionsverringerten, spezifizierten Daten enthaltenen k-nearest-neighbors zu suchen.

19. Das Verfahren nach Anspruch 1, wobei die Daten in einer Datenbank gespeichert sind, und das Verfahren außerdem Schritte enthält, um

eine Dimensionalität der Datenbank zu verringern und die zur Datenbank gehörende Dimensionsverringierungsinformation zu erstellen;

die zur Datenbank gehörende Dimensionsverringierungsinformation zu speichern;

wobei der Partitionierungsschritt auf den Verringerungsschritt reagiert.

20. Das Verfahren nach Anspruch 19, um eine exakte Suche durchzuführen, wobei das Verfahren Schritte enthält, um

die Dimensionalität der spezifizierten Daten basierend auf der Dimensionsverringierungsinformation für die Datenbank zu verringern;

als Reaktion auf die Verringerung die dimensionsverringerten spezifizierten Daten basierend auf der Gruppierungsinformation einem der Cluster zuzuordnen;

eine Dimensionalität der spezifizierten Daten auf diejenige eines dimensionsverringerten Clusters, der von einem zugehörigen Cluster basierend auf der Dimensionsverringierungsinformation für den zugehörigen Cluster definiert wurde, zu verringern; und

basierend auf einer dimensionsverringerten Version der spezifizierten Daten nach einem passenden dimensionsverringerten Cluster zu suchen.

21. Das Verfahren nach Anspruch 19, um eine Ähnlichkeitssuche durchzuführen, wobei das Verfahren Schritte enthält, um

die Dimensionalität der spezifizierten Daten mittels der Dimensionsverringierungsinformationen, die zu der Datenbank gehören, zu verringern;

basierend auf den Gruppierungsinformationen den Cluster zu finden, zu dem die dimensionsverringerten spezifizierten Daten gehören;

basierend auf der Dimensionsverringeringungsinformation für einen identifizierten Cluster die Dimensionalität der dimensionsverringerten spezifizierten Daten zu verringern;

nach einer dimensionsverringerten Version des Clusters zu suchen, zu der die weiter dimensionsverringerten spezifizierten Daten gehören;

über den mehrdimensionalen Index k-Aufzeichnungen wiederzufinden, die den weiter dimensionsverringerten spezifizierten Daten in dem Cluster am ähnlichsten sind; zu beurteilen, ob andere Cluster Aufzeichnungen enthalten können, die näher an den spezifizierten Daten sind als die weitesten der wiedergefundenen k-Aufzeichnungen;

als Reaktion auf den Beurteilungsschritt einen weiteren Cluster zu suchen, der den spezifizierten Daten am nächsten ist; und

die Beurteilung und die Suche für alle anderen Cluster zu wiederholen.

22. Das Verfahren nach Anspruch 19, wobei die in einer Datenbank gespeicherten Daten außerdem Schritte enthalten, um einen oder mehrere dimensionsverringerte suchbare Indizes für einen oder mehrere dimensionsverringerte Cluster zu erstellen und zu speichern.

23. Das Verfahren nach Anspruch 19, um eine exakte Suche durchzuführen, wobei das Verfahren Schritte enthält, um

basierend auf der gespeicherten Gruppierungsinformation die spezifizierten Daten einem der Cluster zuzuordnen;

die spezifizierten Daten in einen dimensionsverringerten Cluster zu zerlegen, der von einem zugehörigen Cluster definiert wurde und die dimensionsverringerten Informationen als Reaktion auf die Zuordnung für den zugehörigen Cluster zu speichern; und

basierend auf den zerlegten spezifizierten Daten diese Indizes nach passenden dimensionsverringerten Clustern zu durchsuchen.

24. Das Verfahren nach Anspruch 23, wobei die Abfrage eine Suchschablone enthält, und das Verfahren außerdem Schritte enthält, wobei

die Zuordnung außerdem die Identifizierung der Suchschablone mit einem Cluster basierend auf der gespeicherten Gruppierungsinformation enthält;

die Zerlegung außerdem die Projektion der Suchschablone in einem Unterraum für einen identifizierten Cluster basierend auf der gespeicherten Dimensionsverringeringungsinformation enthält; und

die Suche außerdem die Durchführung einer internen Cluster-Suche für eine projizierte Schablone enthält.

25. Das Verfahren nach Anspruch 1, das außerdem Schritte enthält, um

(a) um die Cluster-Grenzen, die einer Annäherung 0. Ordnung an die Cluster-Geometrie entsprechen, zu erstellen;

(b) sich mittels eines minimalen Ziehpunktrahmens jedem der Cluster anzunähern und somit eine Annäherung 1. Ordnung an die Geometrie von jedem Cluster zu erstellen;

(c) den Ziehpunktrahmen in $2k$ -Hyperrechtecke zu partitionieren, wobei sich die Partitionierung in einem Mittelpunkt von jeder Dimension befindet;

(d) nur die Hyperrechtecke zurückzubehalten, die Datenpunkte enthalten und somit eine Annäherung 2. Ordnung an die Geometrie des Clusters zu erstellen; und

(e) die Schritte (c) und (d) für jedes der zurückbehaltenen Hyperrechtecke zu wiederholen, um nacheinander

die Annäherung 3., 4. usw. n-ter Ordnung an die Geometrie des Clusters zu erstellen.

26. Das Verfahren nach Anspruch 25, um eine Hierarchie mit Annäherungen an die Geometriestruktur für jeden Cluster zu suchen, wobei das Verfahren außerdem Schritte enthält, um

die Dimensionalität der spezifizierten Daten mittels der zur Datenbank gehörenden Dimensionsverringere-
information zu verringern;

basierend auf der Gruppierungsinformation den Cluster zu finden, zu dem die dimensionsverringerten spezi-
fizierten Daten gehören;

basierend auf der Dimensionsverringereinformation für einen identifizierten Cluster die Dimensionalität
der dimensionsverringerten spezifizierten Daten zu verringern;

nach einer dimensionsverringerten Version des Clusters zu suchen, zu dem die weiter dimensionsverringerten
Daten gehören;

über den mehrdimensionalen Index die k-Aufzeichnungen, die den weiter dimensionsverringerten spezifizier-
ten Daten im Cluster am ähnlichsten sind, wiederzufinden;

zu beurteilen, ob einer oder mehrere andere Cluster Aufzeichnungen enthalten können, die näher an den
spezifizierten Daten sind als die, die von den wiedergefundenen k-Aufzeichnungen am weitesten entfernt sind;

einen anderen Cluster nur dann zurückzubehalten, wenn er einen der k-nearest-neighbors mit den spezifi-
zierten Daten, die auf den Grenzen des Clusters basieren, enthalten kann;

basierend auf den zunehmend feineren Annäherungen an die Geometrie iterativ zu bestimmen, wenn ein
zurückbehaltener Cluster einen der k-nearest-neighbors enthalten könnte, und den zurückbehaltenen Cluster
nur dann zurückzubehalten, wenn der Cluster in der feinsten Ebene in der Hierarchie der aufeinanderfolgenden
Annäherungen akzeptiert wird; und

einen zurückbehaltenen Cluster als einen Kandidaten-Cluster zu identifizieren, der als Reaktion auf den Schritt
der iterativen Bestimmung einen oder mehrere der k-nearest-neighbors mit Daten enthält.

27. Eine maschinenlesbare Programmspeichereinheit, die einen oder mehrere dimensionsverringerte Indizes in mehr-
dimensionalen Daten enthält, wobei die Programmspeichereinheit ein Programm mit Anweisungen ausführen
kann, die von der Maschine ausführbar sind, um Verfahrensschritte durchzuführen und mehrdimensionale Daten
wie in Anspruch 1 angemeldet darzustellen.

28. Ein Computerprogrammprodukt mit:

einem computernutzbaren Medium mit computerlesbaren Programmcodemittel, die hier zur Darstellung mehr-
dimensionaler Daten ausgeführt werden, wobei die computerlesbaren Programmcodemittel in dem Compu-
terprogrammprodukt enthalten:

computerlesbare Programmcodemittel zur Gruppierung, um einen Computer zu veranlassen, die Parti-
tionierung der mehrdimensionalen Daten in einen oder mehrere Cluster durchzuführen;

computerlesbare Programmcodemittel, die mit den Gruppierungsmittel verbunden sind, um einen Com-
puter zu veranlassen, Gruppierungsinformationen für einen oder mehrere Cluster zu erstellen und zu
speichern;

computerlesbare Programmcodemittel zur Dimensionsverringereung, die mit den Gruppierungsmittel ver-
bunden sind, um einen Computer zu veranlassen, einen oder mehrere dimensionsverringerte Cluster und
Dimensionsverringereinformationen für einen oder mehrere Cluster zu erstellen; und

computerlesbare Programmcodemittel, die mit Dimensionsverringereungsmittel verbunden sind, um einen
Computer zu veranlassen, die Dimensionsverringereinformation zu speichern.

Revendications

1. Procédé informatique pour la représentation de données multidimensionnelles, comprenant les phases qui consistent à :

a. Partager les données multidimensionnelles en un ou plusieurs groupes ;
 b. Générer et enregistrer des informations de regroupement pour le(s)dit(s) groupe(s) ;
 c. Générer un ou plusieurs groupes aux dimensions réduites et des informations du degré de réduction pour le(s)dit(s) groupe (s) ; et
 d. Enregistrer les informations relatives au degré de réduction.

2. Procédé selon la revendication 1, comprenant en outre les phases qui consistent à :

générer et enregistrer un index des degrés de réduction des dimensions pour le(s)dit(s) groupe(s) aux dimensions réduites.

3. Procédé selon la revendication 1, où les données sont enregistrées soit dans une base de données spatiales soit dans une base de données multimédia qui comprend une pluralité d'enregistrements de données qui contiennent chacun une pluralité de champs, comprenant en outre les phases suivantes :

créer une représentation de la base de données à indexer sous la forme d'un ensemble de vecteurs, où chaque vecteur correspond à une rangée dans la base de données et les éléments de chaque vecteur correspondent aux valeurs, pour une rangée précise, contenues dans les colonnes pour lesquelles l'index consultable sera généré ; et
 ledit partage comprenant le fait de diviser les vecteurs en le(s)dit(s) groupe(s).

4. Procédé selon la revendication 2, comprenant en outre la phase qui consiste à enregistrer l'index entièrement dans la mémoire principale d'un ordinateur.

5. Procédé selon la revendication 2, où ladite création d'un groupe aux dimensions réduites comprend une décomposition en valeur singulière comprenant en outre les phases qui consistent à :

produire une matrice de transformation et les valeurs propres de la matrice de transformation pour chacun desdits groupes ; et
 sélectionner un sous-ensemble des valeurs propres comprenant les plus grandes valeurs propres ; où les informations de réduction des dimensions comprennent la matrice de transformation et le sous-ensemble des valeurs propres ;

6. Procédé selon la revendication 5, pour rechercher les k enregistrements les plus semblables à des données spécifiées, en utilisant l'index de degrés de réduction des dimensions, comprenant les phases suivantes :

associer des données spécifiées au(x)dit(s) groupe(s), sur la base des informations de regroupement qui ont été enregistrées ;
 projeter les données spécifiées dans un sous-espace d'un groupe associé, en se basant sur les informations de réduction des dimensions relatives au groupe associé ;
 en réponse à ladite projection, générer des données de réduction des dimensions incluant un complément orthogonal pour les données spécifiées qui ont été projetées ;
 rechercher, par le biais de l'index, le groupe associé qui a k enregistrements ressemblant le plus aux données spécifiées projetées ;
 déterminer si un quelconque autre groupe peut comprendre l'un quelconque des k enregistrements les plus semblables aux données spécifiées projetées ;
 répéter ladite recherche d'un dit quelconque groupe qui puisse comprendre l'un quelconque des k enregistrements les plus semblables aux données spécifiées projetées.

7. Procédé selon la revendication 6, où les données spécifiées comprennent un modèle de recherche, comprenant en outre les phases suivantes :

ladite phase de projection incluant la projection du modèle, en utilisant les données de réduction des dimen-

sions, sur un sous-espace associé au groupe auquel il appartient ;
générer des informations de réduction des dimensions du modèle pour un modèle projeté ; où ladite recherche, par le biais de l'index, est basée sur le modèle projeté et sur les données de réduction des dimensions du modèle ; et
mettre à jour un ensemble de k voisins les plus proches des k enregistrements les plus semblables au modèle de recherche.

8. Procédé selon la revendication 5, où ladite méthode de sélection d'un sous-ensemble des valeurs propres est une fonction de précision et de rappel des résultats renvoyés.

9. Procédé selon la revendication 2, pour rechercher k enregistrements ressemblant le plus à des données spécifiées, le procédé de recherche comprenant les phases qui consistent à :

identifier un groupe auquel appartiennent les données spécifiées, en se basant sur les informations de regroupement ;
réduire les dimensions des données spécifiées sur la base des informations de réduction des dimensions relatives à un groupe identifié ;
générer des données de réduction des dimensions pour les données spécifiées dont les dimensions ont été réduites, en réponse à ladite réduction ;
rechercher dans l'index multidimensionnel, à l'aide de l'information de réduction des dimensions, une version à dimensions réduites du groupe auquel les données spécifiées appartiennent ;
récupérer, par le biais de l'index multidimensionnel, les k enregistrements les plus semblables dans le groupe ;
identifier les autres groupes éventuels pouvant contenir des enregistrements qui se rapprochent davantage des données spécifiées que celui qui, parmi les k enregistrements les plus semblables récupérés, en est le plus éloigné ;
rechercher un autre groupe possible qui soit le plus proche des données spécifiées, en réponse à ladite phase de détermination ; et
répéter lesdites phases d'identification et de recherche pour tous lesdits autres groupes.

10. Procédé selon la revendication 6 ou 9, comprenant en outre la phase consistant à :

calculer les distances (D) entre les k voisins les plus proches dans la version du groupe et les données spécifiées projetées sous la forme d'une fonction d'un index d^2 de désadaptation où, d^2 (modèle, élément) = D^2 (modèle_projeté, élément) + $\alpha \sim \text{complément_orthogonal} \sim^2$.

11. Procédé selon la revendication 1, où les informations de regroupement comprennent des données relatives à un centre de gravité du(des)dit(s) groupe(s), comprenant en outre la phase qui associe le centre à une étiquette unique.

12. Procédé selon la revendication 1, où le degré de réduction des données est > 8 .

13. Procédé selon la revendication 1, pour réaliser une recherche d'exactitude, comprenant les phases qui consistent à :

associer des données spécifiées à l'un des groupes, en se basant sur les informations de regroupement qui ont été enregistrées ;
réduire les dimensions des données spécifiées en se basant sur l'information relative à la réduction des dimensions pour une version réduite d'un groupe, en réponse à ladite phase d'association ; et
rechercher, sur la base des données spécifiées aux dimensions réduites, la version réduite du groupe qui correspond aux données spécifiées.

14. Procédé selon la revendication 13, où ladite recherche comprend en outre la phase qui consiste à conduire un examen linéaire pour trouver une concordance avec les données spécifiées.

15. Procédé selon la revendication 1, comprenant en outre les phases suivantes :

créer une hiérarchie de groupes aux dimensions réduites en appliquant de manière récursive lesdites phases a) à d) ; et

générer et enregistrer un ou plusieurs index des faibles dimensions pour les groupes situés dans les bas niveaux de ladite hiérarchie.

16. Procédé selon la revendication 15, pour réaliser une recherche d'exactitude, comprenant les phases suivantes :

appliquer de manière répétée les phases qui consistent à :

trouver un groupe auquel appartiennent les données spécifiées, en utilisant les informations de regroupement enregistrées ; et
réduire les dimensions des données spécifiées en utilisant les informations de réduction des dimensions enregistrées, jusqu'à atteindre un niveau correspondant parmi les niveaux les plus bas de la hiérarchie ; et
rechercher, en utilisant les index des faibles dimensions, une version en dimensions réduites du groupe concordant avec les données spécifiées.

17. Procédé selon la revendication 15, pour effectuer une recherche de similitude, comprenant en outre les phases qui consistent à :

appliquer de manière récursive les phases suivantes :

trouver le groupe auquel appartiennent les données spécifiées, en utilisant les informations de regroupement enregistrées ; et
réduire les dimensions des données spécifiées pour correspondre au niveau le plus faible de la hiérarchie des groupes aux dimensions réduites, en utilisant les informations sur les réductions de dimensions qui ont été enregistrées ;
rechercher les derniers groupes possibles qui peuvent contenir un ou plusieurs des k voisins les plus proches des données spécifiées à chaque niveau de la hiérarchie des groupes aux dimensions réduites, en commençant par un dernier groupe à un niveau faible de ladite hiérarchie auquel les données spécifiées appartiennent ; et
pour chaque dernier groupe possible, effectuer une recherche à l'intérieur du groupe pour trouver les k voisins les plus proches des données spécifiées.

18. Procédé selon la revendication 15, pour exécuter une recherche de similitude, comprenant en outre les phases qui consistent à :

réduire les dimensions des données spécifiées ;
appliquer de manière récursive les phases suivantes : trouver le groupe auquel appartiennent les données spécifiées aux dimensions réduites, en utilisant les informations de regroupement enregistrées ; et réduire les dimensions des données spécifiées aux dimensions réduites pour correspondre à un niveau le plus faible de la hiérarchie des groupes aux dimensions réduites, en utilisant les informations de réduction des dimensions qui ont été enregistrées ;
rechercher les derniers groupes pouvant contenir un ou plusieurs des k voisins les plus proches des données spécifiées aux dimensions réduites à chaque niveau de la hiérarchie des groupes aux dimensions réduites, en commençant par un dernier groupe à un niveau le plus faible de ladite hiérarchie auquel les données spécifiées appartiennent ; et
pour chaque dernier groupe possible, effectuer une recherche à l'intérieur du groupe pour trouver les k voisins les plus proches des données aux dimensions réduites spécifiées.

19. Procédé selon la revendication 1, où les données sont stockées dans une base de données, comprenant en outre les phases qui consistent à :

réduire une dimension de la base de données et générer de l'information de réduction des dimensions associée à la base de données ; et
enregistrer l'information de réduction des dimensions associée à la base de données ;
où ladite phase de partage répond à ladite phase de réduction.

20. Procédé selon la revendication 19, pour exécuter une recherche d'exactitude, comprenant les phases qui consistent à :

réduire les dimensions de données spécifiées, en se basant sur les informations de réduction des dimensions pour la base de données;

associer les données spécifiées aux dimensions réduites à l'un des groupes, sur la base des informations de regroupement, en réponse à ladite réduction ;

réduire les dimensions des données spécifiées en appliquant le degré de réduction d'un groupe aux dimensions réduites défini par un groupe associé, en se basant sur les informations de réduction des dimensions du groupe associé ; et

rechercher un groupe aux dimensions réduites qui concorde, en se basant sur une version aux dimensions réduites des données spécifiées ;

21. Procédé selon la revendication 19, pour effectuer une recherche de similitude, comprenant les phases suivantes :

réduire les dimensions de données spécifiées en utilisant les informations de réduction des dimensions associées à la base de données ;

trouver le groupe auquel appartiennent les données spécifiées aux dimensions réduites, en se basant sur les données de regroupement ;

réduire la dimension des données spécifiées avec dimensions réduites, en se basant sur les informations de réduction des dimensions pour un groupe identifié ;

rechercher une version en dimensions réduites du groupe auquel appartiennent les données spécifiées dont les dimensions ont été réduites deux fois ;

retrouver, par le biais de l'index multidimensionnel, les k enregistrements ressemblant le plus aux données spécifiées dont les dimensions ont été réduites deux fois dans le groupe ;

déterminer si d'autres groupes peuvent contenir des enregistrements plus proches des données spécifiées que le plus éloigné des k enregistrements récupérés ;

rechercher un autre groupe le plus proche des données spécifiées, en réponse à ladite phase de détermination ; et

répéter lesdites détermination et recherche pour tous lesdits autres groupes.

22. Procédé selon la revendication 19, où les données sont stockées dans une base de données, comprenant en outre la phase qui consiste à générer et enregistrer un ou plusieurs index des dimensions réduites consultable(s), pour le(s)dit(s) groupe(s) aux dimensions réduites.

23. Procédé selon la revendication 19, pour effectuer une recherche d'exactitude, comprenant les phases qui consistent à :

associer des données spécifiées à l'un des groupes, sur la base des informations de regroupement enregistrées ;

décomposer les données spécifiées en un groupe aux dimensions réduites, défini par un groupe associé et les informations sur la réduction des dimensions qui ont été enregistrées pour le groupe associé, en réponse à ladite association ; et

rechercher dans lesdits index un groupe aux dimensions réduites qui corresponde, en se basant sur les données spécifiées décomposées.

24. Procédé selon la revendication 23, où la requête comprend un modèle de recherche, comprenant en outre les phases suivantes :

ladite association comprenant d'identifier un groupe au moyen du modèle de recherche à en se basant sur les informations de regroupement enregistrées ;

ladite décomposition comprenant la projection du modèle de recherche sur un sous-espace pour un groupe identifié, en se basant sur les informations de réduction des dimensions enregistrées ; et

ladite recherche comprenant la recherche, à l'intérieur du groupe, d'un modèle projeté.

25. Procédé selon la revendication 1 comprenant en outre les phases qui consistent à :

a. Générer des limites pour les groupes, qui correspondent à une approximation d'ordre zéro d'une géométrie desdits groupes ;

b. Faire l'approximation de la géométrie de chacun des groupes, au moyen d'un cadre limite et générer à partir d là une approximation d'ordre premier de la géométrie de chaque groupe ;

- c. Diviser le cadre-limite en $2k$ hyper rectangles, où ladite division se fait au point central de chaque dimension ;
 d. Ne retenir que les hyper rectangles qui contiennent des points de données et générer à partir d'eux une approximation d'ordre second de la géométrie du groupe
 e. Répéter lesdites phases c) et d) pour chaque hyper rectangle retenu pour générer successivement les approximations d'ordre troisième, quatrième, ... , $n^{\text{ième}}$ de la géométrie du groupe.

26. Procédé selon la revendication 25, pour effectuer une recherche dans une hiérarchie d'approximations de la structure géométrique de chaque groupe, comprenant en outre les phases suivantes :

- réduire les dimensions des données spécifiées en utilisant les informations de réduction des dimensions associées à la base de données ;
 trouver le groupe auquel appartiennent les données spécifiées aux dimensions réduites, en se basant sur les informations de regroupement ;
 réduire les dimensions des données aux dimensions réduites spécifiées, en se basant sur les informations de réduction des dimensions pour un groupe identifié ;
 récupérer, par le biais de l'index multidimensionnel, k enregistrements ressemblant le plus aux données spécifiées aux dimensions réduites deux fois dans le groupe ;
 déterminer si un ou plusieurs autres groupes peuvent contenir des enregistrements plus proches des données spécifiées que celui des k enregistrements récupérés qui en est le plus éloigné ;
 ne retenir un autre groupe que s'il peut contenir l'un quelconque des k plus proches voisins des données spécifiées, en se basant sur les limites du groupe ;
 déterminer de manière itérative si un groupe retenu peut contenir l'un quelconque des k plus proches voisins, en se basant sur des approximations de plus en plus fines de la structure géométrique, et ne retenir le groupe retenu que si le groupe est accepté au niveau le plus fin de la hiérarchie des approximations successives ; et
 identifier un groupe retenu comme étant un groupe possible contenant un ou plusieurs des k plus proches voisins des données, en réponse à ladite phase de détermination itérative.

27. Dispositif de stockage de programmes lisible par machine, qui comprend un ou plusieurs index aux dimensions réduites relatifs à des données multidimensionnelles, le dispositif de stockage de programme mettant en oeuvre tangiblement un programme d'instructions pouvant être exécuté par la machine pour exécuter les phases d'un procédé de représentation de données multidimensionnelles selon la revendication 1.

28. Programme informatique comprenant :

- un support utilisable par un ordinateur, dans lequel est mis en oeuvre un moyen d'un code de programme lisible par machine qui représente des données multidimensionnelles, le moyen du code de programme lisible par machine dans ledit programme informatique comprenant :

- un moyen de regroupement du code de programme lisible par machine, par lequel un ordinateur partage les données multidimensionnelles en un ou plusieurs groupes ;
 un moyen du code de programme lisible par machine, couplé audit moyen de regroupement, par lequel un ordinateur génère et stocke des informations de regroupement pour le (s) dit (s) groupe(s) ;
 un moyen de réduction des dimensions du code de programme lisible par machine, couplé audit moyen de regroupement, par lequel un ordinateur génère un ou plusieurs groupes aux dimensions réduites et des données de réduction des dimensions pour le(s)dit(s) groupe(s) ; et
 un moyen du code de programme lisible par machine, couplé audit moyen de réduction des dimensions, par lequel un ordinateur enregistre les informations de réduction des dimensions.

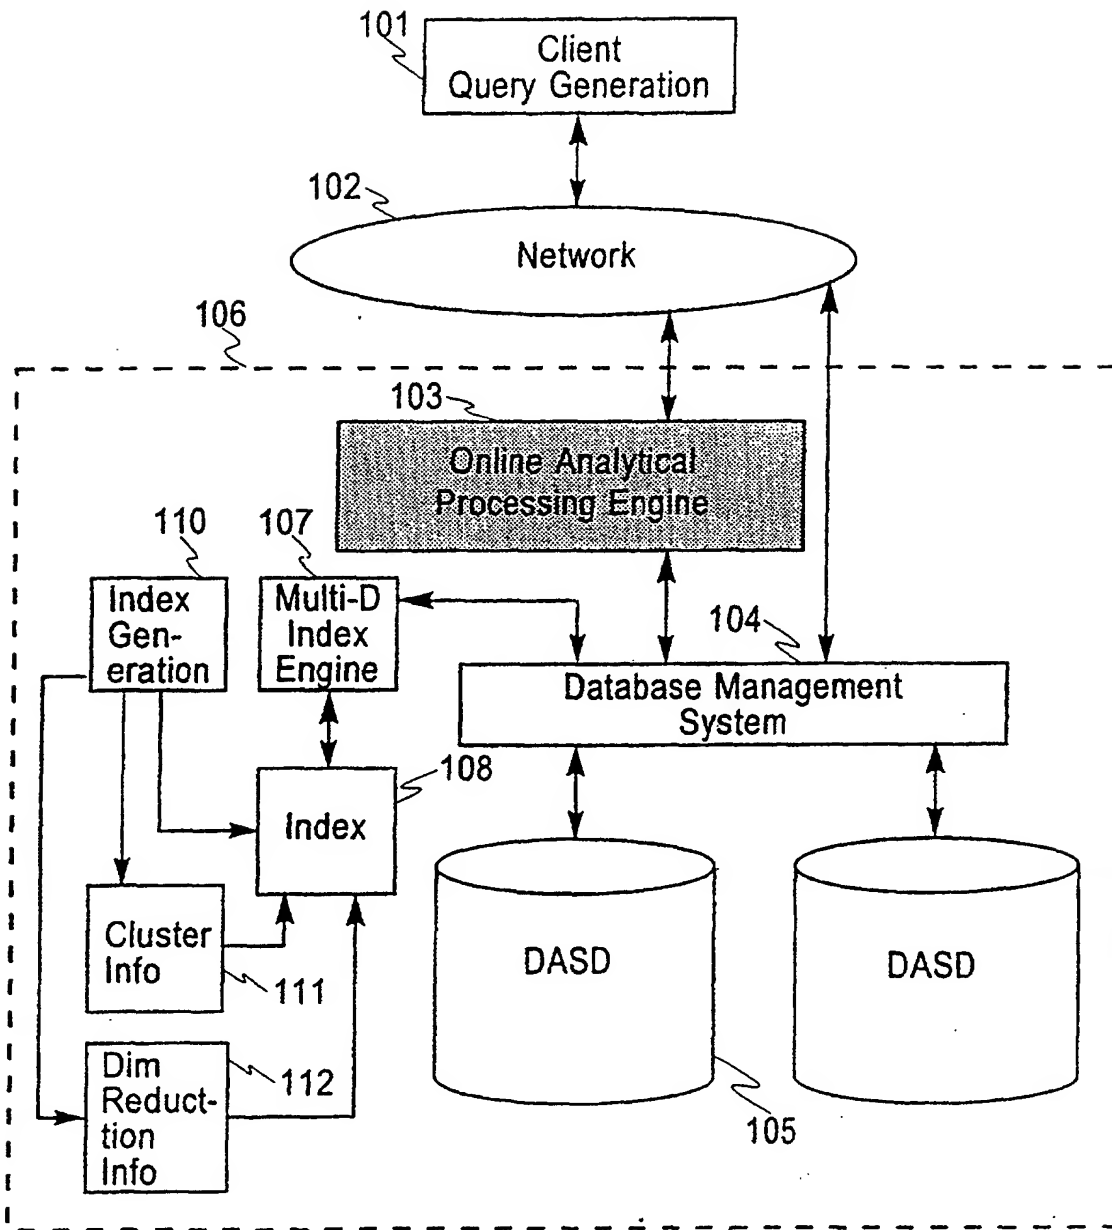


Fig. 1




Fig. 2




Fig. 3




Fig. 4.


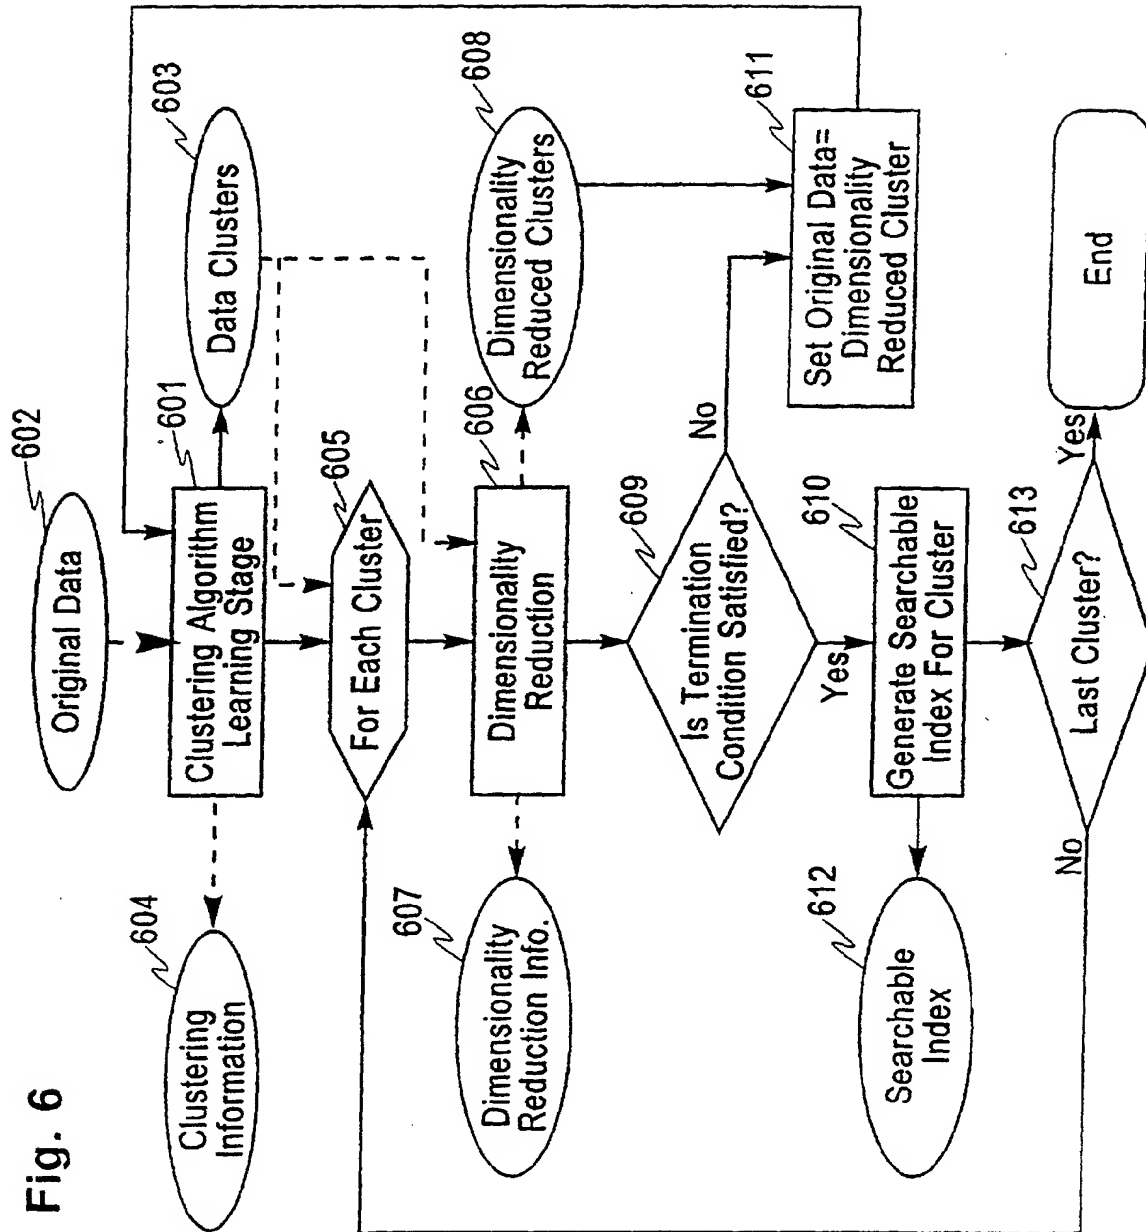


Fig. 5

Fig. 6



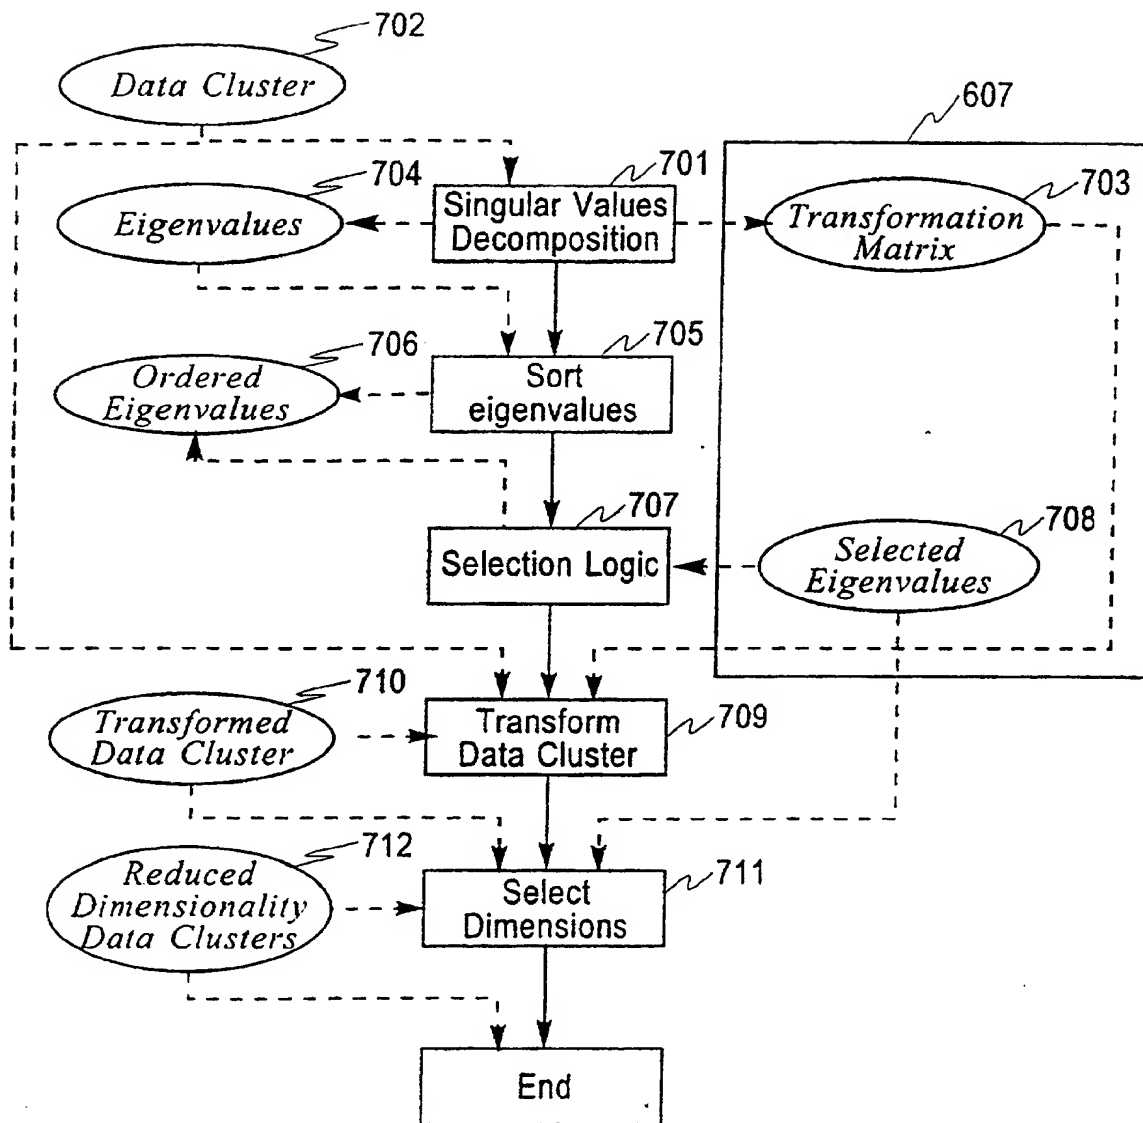


Fig. 7

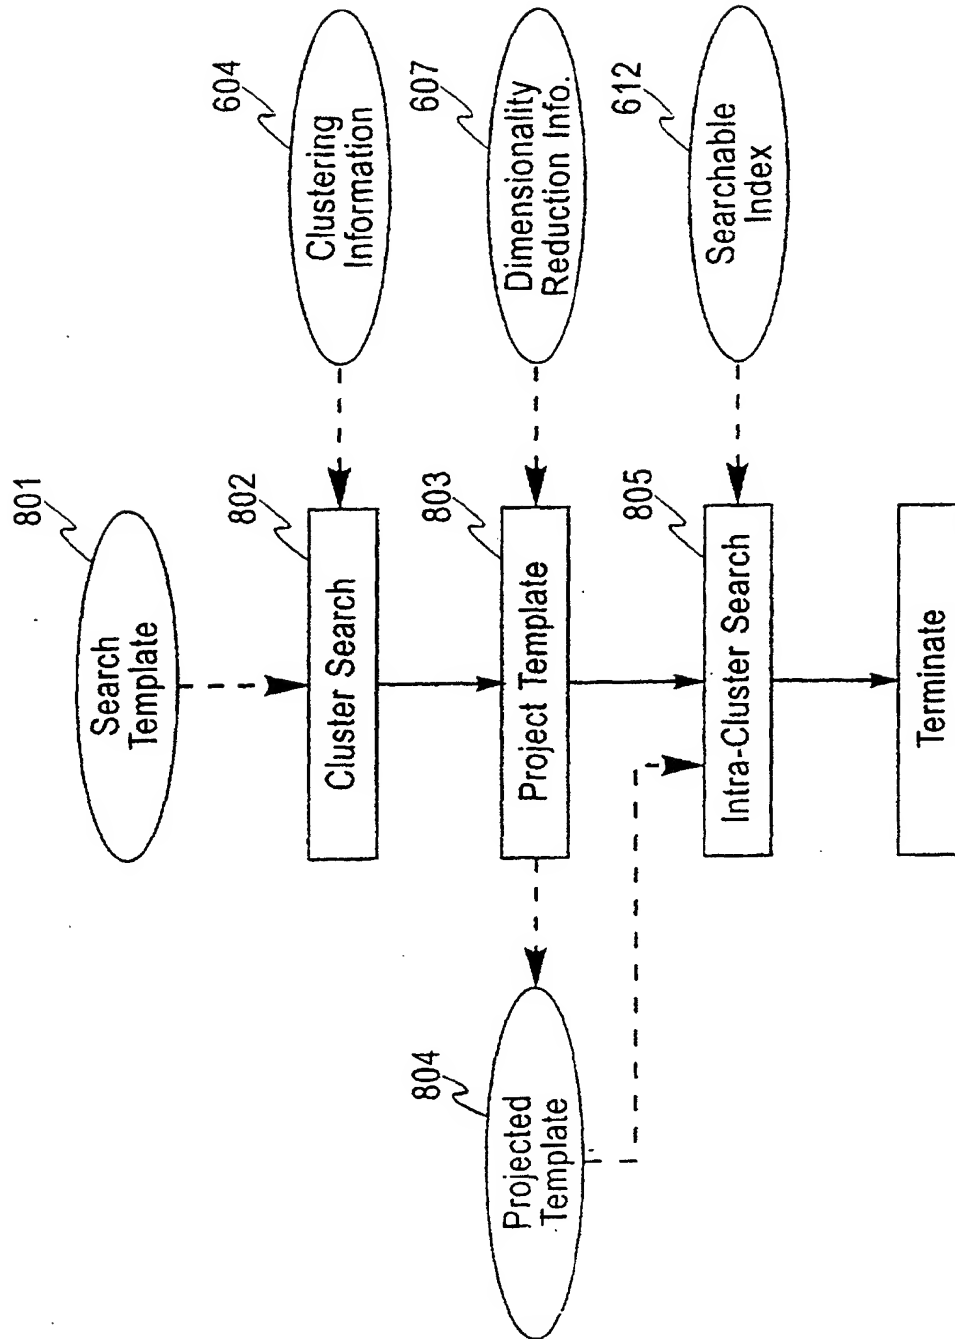


Fig. 8

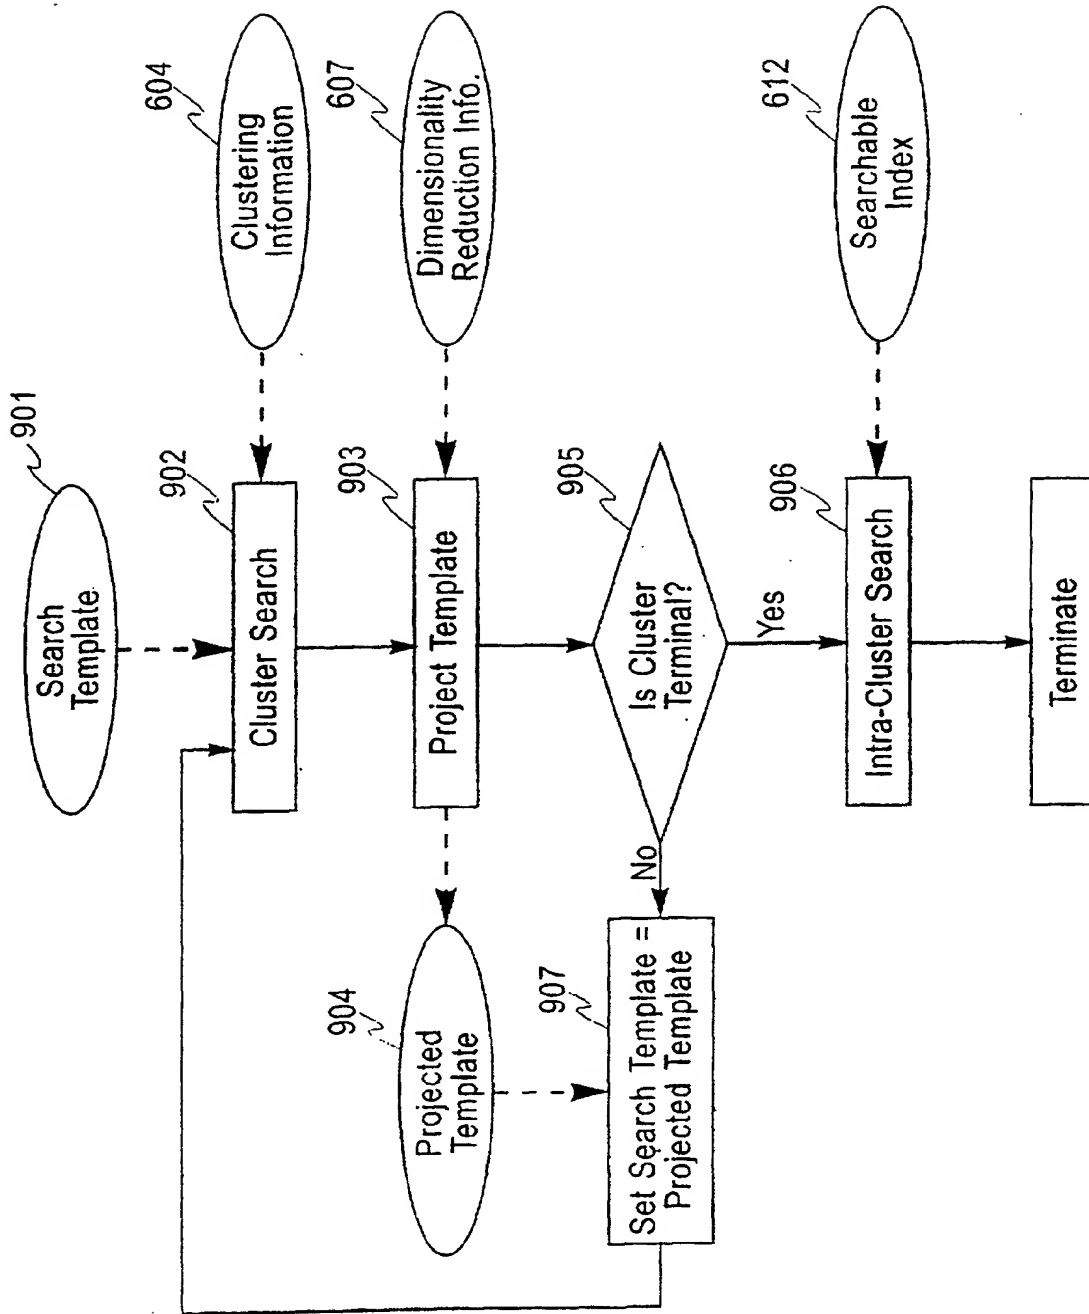


Fig. 9

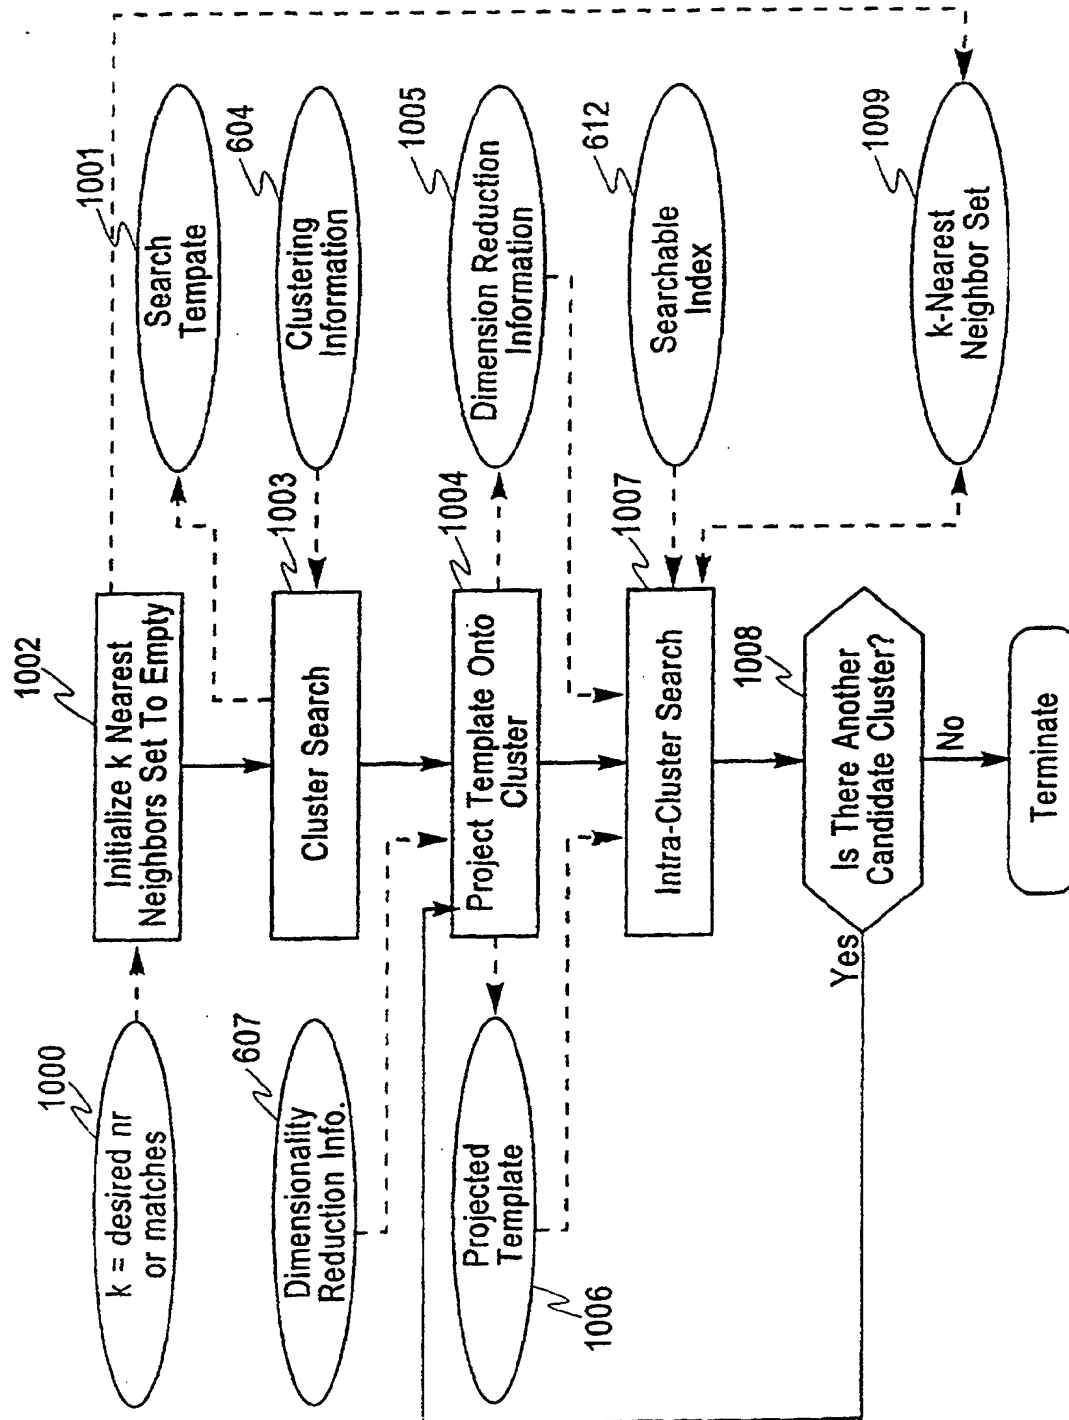


Fig. 10

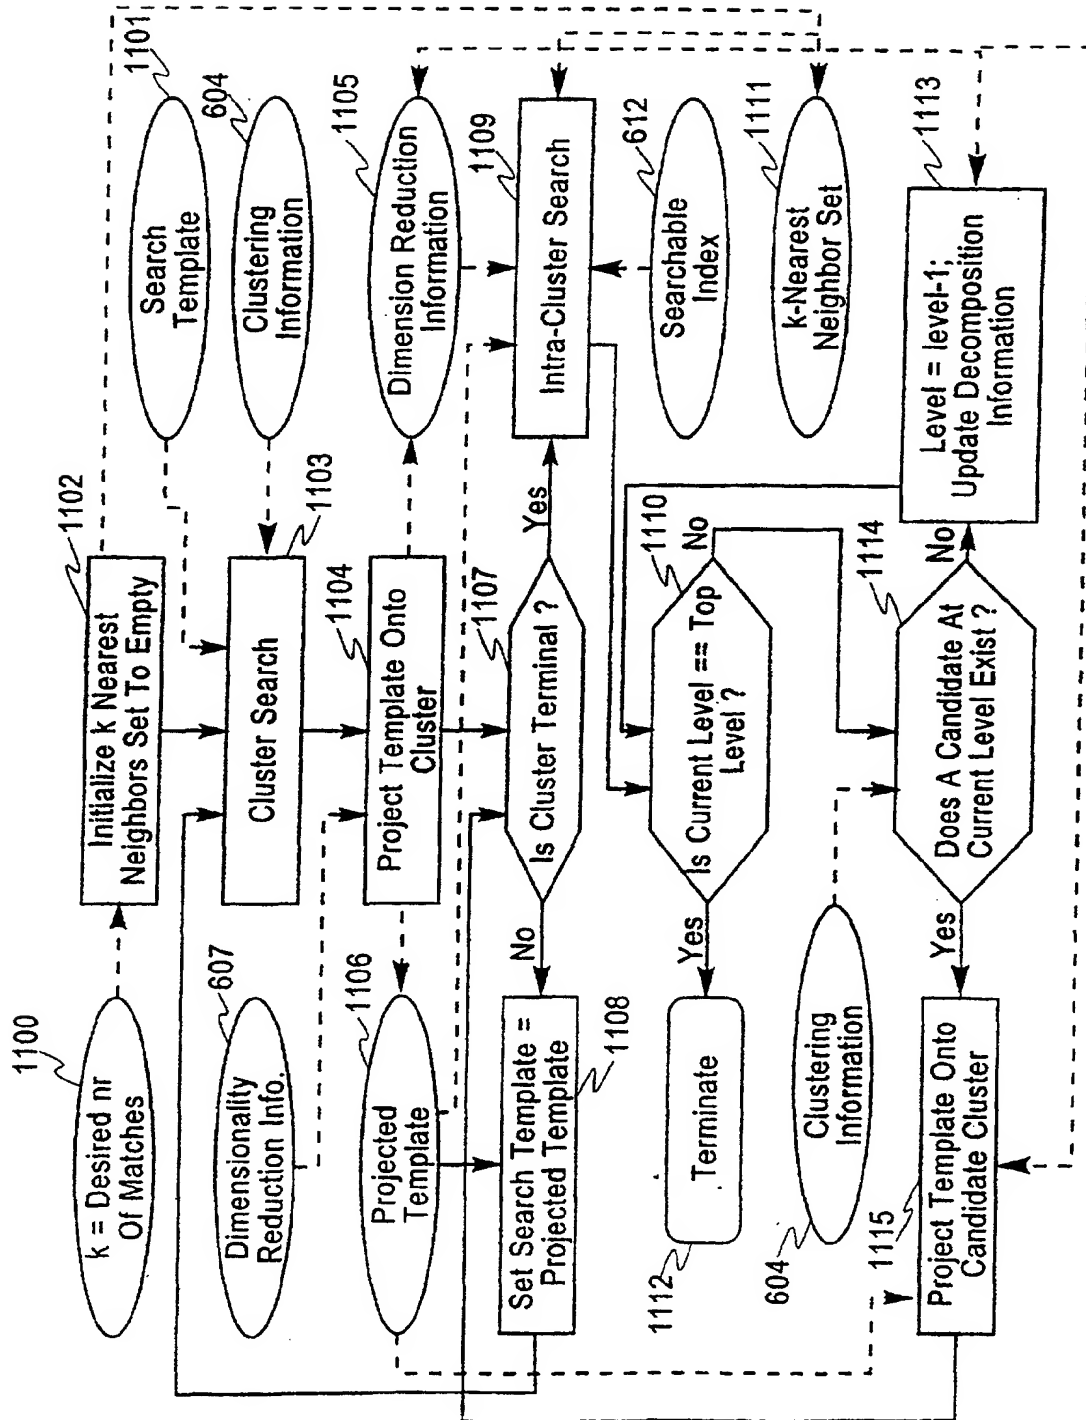



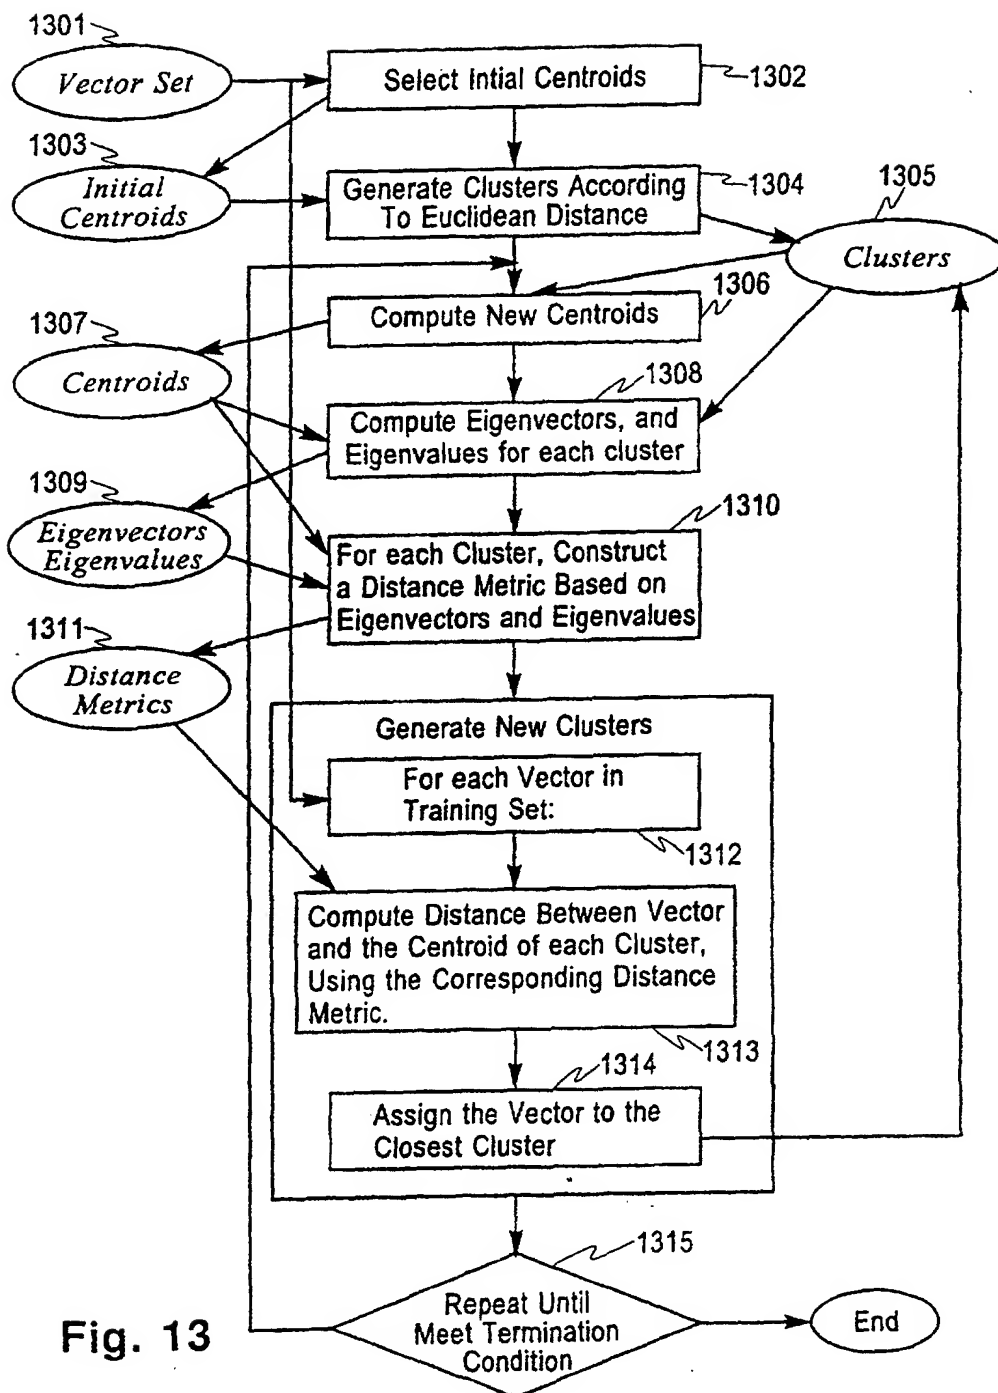


Fig. 11








Fig. 14

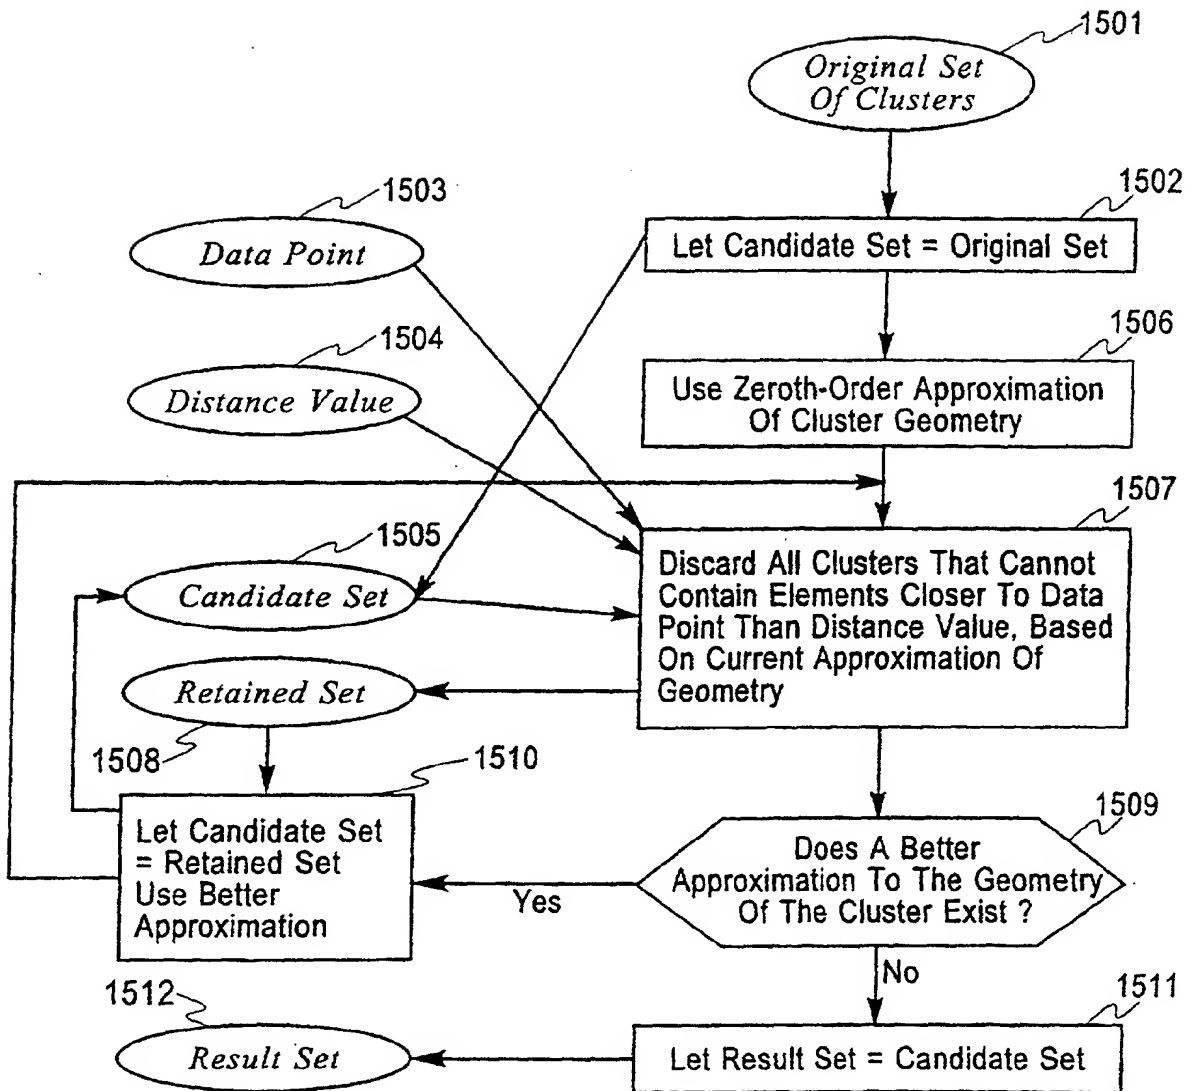


Fig. 15